The Impact of Lead Time on Truckload Transportation Rates

by

Erik R. Caldwell

B.S. History
United States Military Academy, 1996

and

Bryan C. Fisher

M.S. Computer Science
The Johns Hopkins University, 2005

B.S. Computer Science
Brigham Young University, 2001

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of

Master of Engineering in Logistics

at the

Massachusetts Institute of Technology

June 2008

© 2008 Erik R. Caldwell and Bryan C. Fisher. All rights reserved.

The authors hereby grant to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Signatures of Authors...
Master of Engineering in Logistics Program, Engineering Systems Division
May 19, 2008

Certified by.......... ....
Dr. Chris Caplice
Executive Director, Masters of Engineering in Logistics Program
Executive Director, Center for Transportation and Logistics
Thesis Supervisor

Accepted by...
Prof. Yossi Sheffi
Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, Center for Transportation and Logistics
Director, Engineering Systems Division

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
AUG 06 2008
LIBRARIES

ARCHIVES

The Impact of Lead Time on Truckload Transportation Rates

by

Erik R. Caldwell and Bryan C. Fisher

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of
Master of Engineering in Logistics

Abstract

The objective of this thesis was to analyze truckload shipment transactions in order to determine if rates are impacted by tender lead time, which is the amount of time between when a carrier is offered a load to when the load needs to be picked up. The research specifically focused on how tender rejections by carriers are the ultimate driver of transportation cost variances since most rates are contractually fixed in advance. The data revealed a strong correlation between tender rejections and increased costs. Many factors affect transportation costs. The transportation model in the paper included three key baseline factors: distance, origin, and destination of the load. The model also included tender and pick up day of week activity, economies of scale, carrier size, and tender lead time to quantify how the factors influence the cost of a load. The research suggests that even though the baseline factors dominate the cost of most loads, shippers can create savings by modifying business policy with regard to tender lead time and other factors included in the model.

Thesis Supervisor: Dr. Chris Caplice
Title: Executive Director, Center for Transportation and Logistics

Acknowledgements

My portion of the paper is dedicated to my grandfather, Bert Russell (1922-2008), who spent his life studying and executing the movement of goods and people. He demonstrated his expertise in this field as a pilot in the Army Air Corps serving in one of the most dangerous jobs in WWII as a pilot flying "Over the Hump" (Himalayan Mountains) from India into China, as a Captain at the developing national airline carrier Delta, and in his passion for railroads. He also helped five children and eleven grandchildren as they transited from one stage of life to the next through his example and counsel. Thanks.

-Erik

I would like to thank my wife for her encouragement and support of my academic and professional endeavors. I thank my daughters for their enthusiasm and goodness. I also thank my parents for teaching me to love learning.

Lastly, I acknowledge our advisor, our sponsors, and others who have generously assisted us. Thank you.

-Bryan

TABLE OF CONTENTS

List of Figures

List of Tables

1 Introduction

The US truckload transportation market accounts for over $150B in annual revenue with retail and manufacturing companies typically spending 3-4% of their revenue on transportation (Morgan Stanley 2008). This creates a situation where even small improvements in business policies can drive substantial savings. The purpose of this research is to determine if business policies relating to tender lead time have an impact on transportation costs. Tender lead time refers to the time from when a carrier is offered a load to when the load must be picked up. We found that policies associated with lead time, corridor volume, tender and pick up day of week, and carrier size have a significant impact on the total transportation cost for companies. On average these variables account for roughly 1% of the cost of the load; however, when we summarized the extreme parameters of each variable the savings equaled $287, or 28% of the cost of a typical load. As part of that total, lead time alone could influence the cost of a load by $71 at the extremes of tendering a load with lead time of under one day or over twelve days.

1.1 Motivation

Does the amount of time between when a load is tendered and when it ships affect the overall cost? This would almost seem like an obvious question but when we polled a number of industry veterans we received widely different answers on whether lead time had any impact, the magnitude of the impact, and the belief that lead time matters greatly under certain market conditions and not at all in other market conditions. From our initial research we found that there was an absence of academic or empirical data to support or refute the hypothesis that lead time has an impact in transportation. Lacking insight within the transportation industry we reviewed an activity that highlighted variability in pricing due to lead time - ticket scalping. Like airline flights or hotel rooms, ticket scalping provides a service that has a defined expiration time and a price that changes based on the lead time before an event. From first hand experience outside the Boston Fleet Center we were able to learn that scalped ticket prices follow a yield management curve where scalpers are hesitant to sell tickets too far in advance before the particular game since they could potentially earn more profit closer to the start time. The best time for the consumer was during the national anthem, but still before the game began, and a skilled negotiator was able to buy rink side seats for less than half of face-value price and not miss the puck drop. Interestingly, not long after the game began any remaining tickets became steeply discounted, practically by the minute, and the secondary ticket market would quickly disappear. A similar event happens in the transportation market when carriers fail to secure loads and have idle tractor capacity.

From these observations of a common yield management pricing practice, we were able to draw some very relevant insights that helped guide our research for lead time in transportation management. Although lead time was extremely important in ticket scalping, lead time was not the most dominant variable. The most important variable for the scalpers was whether the game was sold out. If consumers had the ability to still purchase tickets at the counter for face-value then the scalpers were forced to discount their tickets to make them more attractive or risk holding tickets without a secondary ticket market. Other variable factors such as weather, location of the game, home team, away team, and significance of the match up had an effect on overall ticket prices.

Like ticket pricing, transportation pricing was also dominated by a key variable, mileage, which predicted the majority of the cost of a load. The other key factors were origin, destination, lead time, corridor volume, tender day of week, pick up day of week, and carrier size preference. With a better appreciation of pricing impacts we were able to examine the impact of lead time on transportation costs using a more advanced model that incorporated all relevant variables.

The ability to improve transportation costs through an increased understanding of the impact of business policies is particularly important to our research sponsor company, C. H. Robinson. As the largest public transportation and logistics company in the US, C. H. Robinson manages both transportation execution and planning for their clients. All the transactional data used in the research was part of C.H. Robinson's Transportation Management Center (TMC), which provides outsourced transportation management for specific clients and is separate from their traditional brokerage service. The TMC uses a centralized transportation management system for all clients to help plan, manage, and track transportation activities. This centralized aspect of the TMC helped ensure standardization of data. Another important aspect of the TMC is that they leverage best practices in transportation consistently across all customers and carriers. This consistency of planning and execution makes the research much more powerful since it was easier to identify different individual customer behaviors and policies with regard to lead time that are not the result of different approaches to transportation management.

Our research should have broad applicability to other companies that either manage freight or ship products. Shippers can use the research to improve their understanding of what factors result in a higher rejection of load tenders. When load tenders are rejected by the preferred carrier, shippers are forced to resort to more costly carriers on a specific lane, which we define as a pair of five-digit origin and destination zip codes. Companies will be able to review their business policies to calculate the impact their decisions have on the key variables that influence transportation costs. Using this improved knowledge of what affects the cost of loads, the companies can reevaluate their policies with a truer understanding of the overall impact. Finally, the research concludes with a number of suggestions on how to best take advantage of the research that are both practical and relatively easily to implement.

1.2 Current Practices

Most companies who ship products in truckload quantities rely on contract carriers. For each origin-destination combination the shipper maintains a routing guide, or list of carriers, who have agreed to haul loads on the route for an agreed upon price. The shipper ranks the list of carriers in the order of preference. Since cost is a key consideration when contracting loads to carriers, most often rates increase with routing guide depth.

Lane: Memphis - Boston			
Sequence	Carrier Name	Carrier Size	Cost
1	Carrier A	Medium	$900
2	Carrier B	Extra Large	$950
3	Carrier C	Medium	$1,200
4	Carrier D	Extra Large	$1,225
5	Carrier E	Small	$1,500

Table 1.1: Example of hypothetical routing guide for a lane

Each time the shipper needs to ship a load, the shipper consults its routing guide for the specific lane and sends a tender offer to the first carrier either electronically or manually. Once the carrier receives the tender they can either accept or deny the tender. The rate has been contractually agreed upon, but it is understood that a carrier will not always have the assets available to accept the load. If the load is rejected, the shipper tries the next carrier in its routing guide until the load is successfully tendered. Sometimes a dozen tenders or more are required before a carrier is found who will haul the load.

Cycle time is the overall time it takes from when a load enters the system until it is actually picked up by the carrier. As shown in Figure 1.1, there are a number of key measurements within the total time. The first is the total lead time which is the time from when a load enters the system until its pick up by date. The tender lead time measures how long a carrier has from when the load is offered to them until they actually pick up the load. Finally, the carrier's performance can be measured by how timely they actually pick up loads compared to when they were suppose to pick the load up.

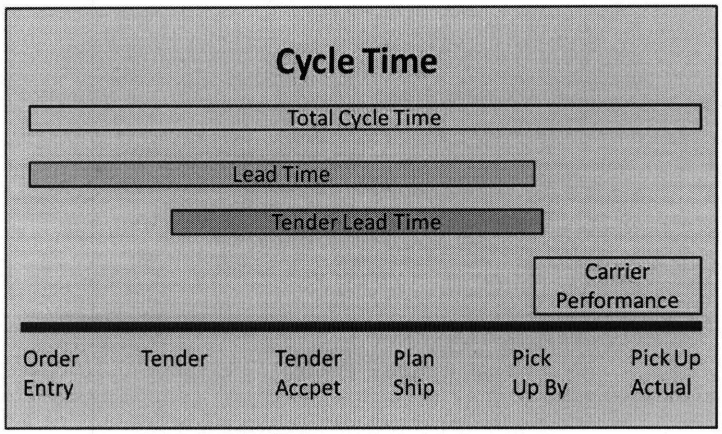

Figure 1.1: Measurements of cycle time

The cycle time can be seen in the process that C. H. Robinson follows to manage their loads. Almost all steps in the process are automated using electronic data interface (EDI) transmissions between the customers' operating system, C.H. Robinson's transportation management system, and the carriers' capacity management system as seen in Table 1.2.

Lifecycle Stage	Process
Load enters TMC system	Automated EDI from Client
Origin pickup time set	Automated EDI from Client
Destination appointment set	*Manual Entry / Phone Call*
Selection of carrier in routing guide	Automated in Transportation System
Tender request sent to carrier	Automated EDI to Carrier
Carrier "accept" / "reject"	Automated EDI from Carrier
→If "reject" then next carrier selected	Automated in Transportation System
Actual pick up time set	Automated EDI from Carrier
Actual arrival at destination	Automated EDI from Carrier
Carrier invoicing	Automated EDI Carrier
Carrier Payment	Automated EDI from TMC

Table 1.2: Lifecycle stages

This highly automated process provides a high degree of reliable data and ensures that most steps happen without need for operator intervention. The one exception is the requirement to set the destination appointment manually (highlighted above) after calling the customer, and delays in this step of the process could be targeted for process improvement to reduce lead time.

In this paper we have limited our research to long-haul dry van shipments which we defined as full truckload shipments of over 250 miles. There are two main reasons for this constraint on the analysis. First, the quantity of long-haul dry van shipments, in terms of the number of loads and also total spend, is by far greater than that of other modes. Second, if an in depth analysis of this most broadly used category reveals a correlation between tender lead time and long-haul dry van rates, then the same techniques could easily be applied to other types of truckload shipments such as refrigerated, flatbed, and intermodal to understand how lead time affects those categories.

1.3 Literature Review

We conducted a review of the literature related to our topic by surveying over 60 transportation and procurement periodicals dating as far back as the 1960s. We also interviewed transportation experts and reviewed industry surveys, past theses, books, and online databases. This section explains the scope of our review and how the research in the field compares to our project.

We did not find any research specifically about tender lead times, but we did examine work done in the area of product lead times, or the interval between when a product is ordered and when it is received. Estimating the length of this interval is one of the inputs for inventory planning equations used in several industries including manufacturing and retail. Our focus, however, is on tender lead time which is independent of product lead time as shown below.

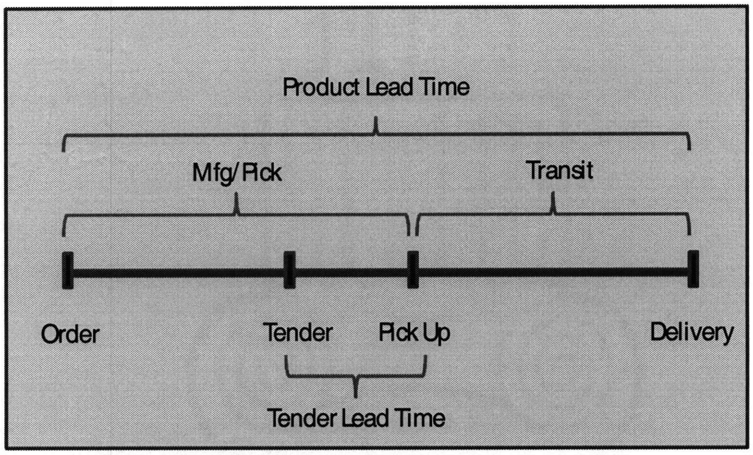

Figure 1.2: Product lead time vs. tender lead time

The freight being shipped should arrive at the same time regardless of how much tender lead time is given. Our hypothesis, however, is that the transportation cost will vary with the tender lead time.

We may be able to draw comparisons with time-sensitive pricing in other industries. Harris and Peacock (1995) have described how several industries change the prices they charge customers based on the amount of time between purchase and when the service is offered. The authors cite the example of Marriott Corporation, which introduced differentiated prices for their hotels in order to increase yield. The resulting revenue allowed the hotel chain to set rates below the previous levels for both those who booked well in advanced and also those who booked at the last minute.

The airline industry also sets prices based in part on the purchase lead time of the reservation. This dynamic pricing practice is part of a larger yield management strategy first introduced by American Airlines and now in widespread use. In this model, service providers change their prices with regard to time and other variables in order to optimize revenue. For example, Anderson and Wilson (2003) report that typically the price for an airline ticket changes based on when it is purchased as shown in Figure 1.3. From the price curve one notices that there is an optimal time when the buyer should purchase a ticket in order to minimize costs.

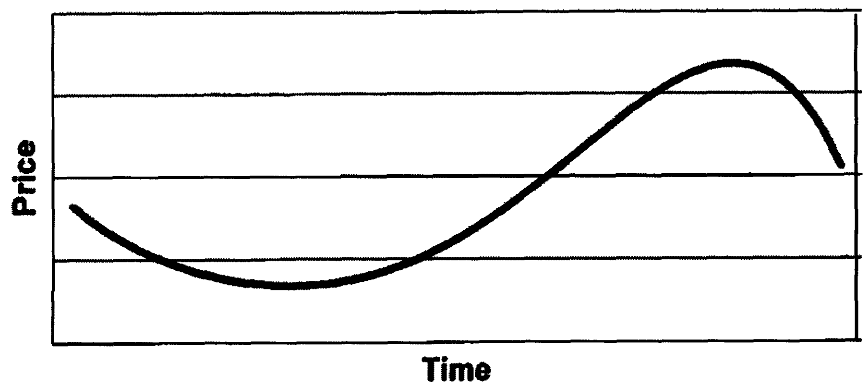

Figure 1.3: Price of airline tickets based on time before flight (Anderson & Wilson 2003)

We are interested in what the shape of this type of graph looks like in the transportation market. There are however some important differences between the truckload transportation industry and industries employing yield management. First, most truckload rates are agreed on contractually between the shipper and the carrier. Once the contract is signed, the carrier's only options are to either accept or decline the load at the contracted rate. However, when shippers do not have a contracted carrier available, or chose not to use a contracted carrier, they can use the spot market to find a carrier and negotiate a one-time price. Second, the airline industry is a consolidated service whereas truckload transportation is a direct service. Since the airline has committed to a schedule, the plane is going to leave even if there are unused seats. Truckload carriers may also have unused capacity in the form of empty trailers and idle drivers. However, trailers can be left in place until a load is located nearby and drivers are generally not paid for the time they spend waiting in between loads. Notwithstanding the differences, some of the effects of yield management may be present in the truckload market.

Swenseth and Godfrey (2002) have shown how expected transportation rates affect economic order quantity (EOQ) calculations. Their research focused on how transportation costs can account for 50% of total logistics spend and as a result have a tremendous impact on all supply chain policy decisions. In particular they studied how to incorporate the effect of variable pricing by the weight of freight shipped. Our study could help build upon their principles by showing that there are other factors that could measurably affect supply chain planning that have previously not been recognized. By incorporating the impact of lead time it would then be possible to calculate the variation by lane with regard to depth of routing guide and expected average rate per load based to a more detailed degree by incorporating average days of lead time. Swenseth and Godfrey demonstrate the importance of creating accurate estimates of freight lane rates rather than just assuming that the best rate is correct. They support this argument citing General Motor's TRANSPART model where it was possible to reduce total annual logistics spend by 26% in one division of the company after incorporating more accurate transportation rate information and its impact on recalculating the economic order quantity.

In reviewing the literature, we found work that did not focus on lead time but more on market thickness, or the amount of demand for a limited carrier base. In particular, Hubbard (2001) found that doubling the demand caused a 30% increase in the number of transactions that occurred outside of routing guide pricing by using the spot market. The sensitivity of market demand leads to both the shippers and carriers attempting to optimize their own value by carriers turning down loads in search of greater profits when there is limited trucking capacity while shippers will more frequently search for cheaper alternatives to their contractual rates in the routing guide when carriers have excess capacity. Hubbard also identified a key link in behavior that shows that carriers only use contracts for longer haul loads and focus on mitigating the risk of overcapacity on specific lanes. Based on this research it is important that we expand this effort to better understand how pricing is affected not just by lead time but also in the broader context of market contractions and expansions with regard to carrier capacity for high volume lanes.

Another alternative to strategically reducing the impact to increased lead time pricing is to use a private fleet. Our research will help support work completed earlier by Mulqueen (2006) that focused on transportation policy for networks that use both contract carriers as well as private fleets. His research helped identify that the most important variables when deciding whether to privatize lanes are volume of freight (corridor volume) and the variability of loads on a lane. Although recalculating Mulqueen's findings after breaking out the lead time cost component would be impractical for the scope of this paper, his work does help us to provide an understanding of how the factors we are researching could reinforce or negate some of his observations from an intuitive approach.

Even though our literature search did not provide direct research related to the impact of lead time within the transportation framework we feel that many other industries tackle a similar challenge through the use of yield management. Our review covered how this occurs in the hospitality and airline industries. Furthermore, the literature helped establish a framework for the importance of transportation pricing, both as an independent cost as well as its affect on total supply chain decisions, helped provide insight into our analysis.

2 Methodology and Analysis of Data

Our analysis can be broken down into four steps: data preparation, data profiling, model building, and model evaluation.

2.1 Data Preparation

The first step in preparing the data for analysis was to cleanse the data of any records with invalid or missing values. Eliminating these transactions was necessary so that the models would be able to operate on the data and produce valid results. In addition to removing erroneous data, we also wanted to make sure that each record represented a long-haul dry van movement, which as mentioned in the introduction is the scope of our analysis. Truckload modes of transport include refrigerated, flatbed, and container; however, dry van shipments vastly outnumber the other modes. We defined long-haul as any movement of at least 250 miles. Since movements shorter than this have less uniform rates and were a smaller percentage of the overall dataset, they were excluded from the analysis. We also ignored loads with extremely low or high per mile rates, assuming that they were recorded incorrectly.

Specifically, we excluded records with any of the following attributes:

- Origin or destination outside of the contiguous 48 states
- Blank tender or pick up date fields
- Blank rate or mileage fields
- Rate per mile greater than $8.00
- Rate per mile less than $.50
- Mode other than dry van
- Miles less than 250
- Customers with less than 200 loads per year
- Any carriers with fixed or contractual rates without regard to tender lead time

Selecting only long haul, dry-van shipments in the contiguous 48 states reduced the dataset from 1,545,053 tender records to 1,050,709. Using the exclusion criteria above further reduced the number of tender records to 521,855. This included the elimination of records that were missing critical fields, unrealistic rates, five customers that did not meet the minimum shipping requirements to be included in the study, and one carrier that had a contractual per mile rate that was not influenced by origin, destination, or lead time. Next, we separated out the accepted tenders from the dataset, resulting in 330,116 distinct loads. Finally, we divided the accepted tenders by year. We used the 273,696 tenders from 2007 as the training dataset to build the models, and we used the remaining 54,400 tenders from 2008 to evaluate the models' predictive performance.

The training dataset from the year 2007 had an average cost per load of $1019 and cost per mile of $1.56. It represented a total transportation spend of over $279M not including accessorial charges.

Each record included a tender sequence number, which in most instances corresponds directly to routing guide depth. 78.3% of loads displayed a sequence 1, meaning that the first carrier that received the tender accepted the load; however, there were loads with tender sequence numbers in the 20s, which means they had been rejected over 19 times before being accepted.

Many of the models we employed do not handle categorical variables, so any categorical variable that we wished to use as an input to these models was decomposed into a series of binary dummy variables, for which a value of one represented TRUE and a value of zero represented FALSE. For example, the origin and destination state variables for each load were recoded into a series of binary variables with the suffix *From-* or *To-*, followed by a two letter state abbreviation.

We also added derived fields to be used during the analysis, based on tender lead time, day of week of tender, day of week of pick up, carrier size, corridor volume and variability, and the Morgan Stanley freight index. We named the tender day of week variables *TenderMon*, *TenderTue*, and so on, and the pick up day of week variables the same way but with the prefix *Pickup-*. We used the annual Transport Topics rankings (Bearth 2007) to create carrier size variables named *CarrierXL* and *CarrierLG*. In our study we consider a corridor to be a pair of three digit zip codes corresponding to the origin and destination of the load. The volume and variability for each customer for each lane were estimated using the data provided. These estimates were coded into the variables CorridorTenderVolume and CorridorTenderWeeks as described in Section 3. Finally, a *FreightIndex* variable was added to each record representing the average Morgan Stanley dry-van freight index (Morgan Stanley 2008) for the week of the tender. This index is a ratio of the estimated demand and the available supply on any given day.

2.2 Data Profiling

After cleansing the data, we employed statistical and manual techniques to profile the data. This step verified that the data appeared to be a representative sample. For example, we made sure that our dataset was not dominated by loads belonging to a single customer or located in a single geographic area. Profiling also allowed us to familiarize ourselves with the data so that we could make better decisions in the model building stage. For example, we calculated statistical correlations between variables so that we could avoid problems arising from multicollinearity among inputs.

2.3 Model Building

We used two approaches to understand the impact of tender lead time. The first approach was to predict how many tenders would be required before a load is successfully tendered, and then estimate how a rate at this level of the routing guide would compare with the rate of the initial tender. The second approach was to predict the rate of the accepted tender directly.

We employed several standard statistical and data mining techniques in our modeling. The most effective model at predicting tenders per load was constructed using the *k*-nearest neighbors technique. This model calculates the distance between each of the training partition records and the record being classified and then makes a prediction based on the output values of the *k* nearest neighbors.

We used least-squares multiple linear regression in both modeling approaches. Linear regression defines a linear relationship between an independent variable, y, and a set of dependent variables, $x_1, x_2, \ldots, x_n$ as follows (Shmueli, Nitin, & Bruce 2007):

$$y = B_0 + B_1x_1 + B_2x_2 + \ldots + B_nx_n + e$$

The set of coefficients B are assigned so that the sum of the squared errors between each observation and the regression line is minimized. B_0 is a constant term and e is the error term. The method has the advantage that each input's contribution to the resulting prediction is simply the product of the input and the coefficient, making it easy to determine how much each of the independent variables is contributing to the prediction.

2.4 Model Evaluation

To aid in evaluating each model the records were partitioned. The largest partition, known as the training partition, was used to construct the model. The remaining records were assigned to the test partition and used to evaluate how well the model performed on new data.

We measured the ability of the multiple linear regression models to explain the variability of the training data using adjusted R^2, defined as follows (Shmueli, et al. 2007):

$$Adjusted\ R^2 = 1 - [(n - 1) / (n - p - 1)] * (1 - R^2)$$

In the equation, p is the number of predictors used in the model and R^2 is the squared correlation coefficient. The resulting adjusted R^2 is a number between zero and one that reflects how much of the variability of the data is being explained by the model, with a penalty for using more inputs. The closer the R^2 value is to 1 the more accurate the regression is with respect to the dataset.

The predictive ability of the models was evaluated against the training dataset, using measures such as error rate and the sum of the squared residuals. These measures helped guide the parameters used for subsequent iterations of the model. Models that perform very well or even perfectly on a training partition are not likely to perform as well on new data. When this happens, the model is said to have been overfit to the data in the training partition. To avoid overfitting we used the previously unseen data in the test partition to estimate how the model would perform generally. By iteratively building and evaluating variations of each type of model, we were able to select model parameters that were best suited for this specific application.

2.5 Dataset Profile

This overview portion of the analysis section is designed to describe what the dataset generally looks like with regard to cycle time pricing as well as to identify areas of opportunity to focus the detailed models. This section takes an approach at a higher level in order to see broad impacts and correlations in the data and to also help the reader understand where our research started so they can follow the logic of which variables stood out as having an impact on transportation pricing. Our later models will provide additional quantitative analysis of the dataset with an emphasis on isolating these different variables to better understand how they interact to affect transportation pricing.

The data provided by the sponsor proved to be very robust in terms of both volume and geographic representation. The map below shows how the dataset covers all of the regions of the contiguous 48 states. The percentage listed above each region's label identifies the percentage of the 334,113 total loads in the dataset shipped from that region. The percentage below each region's label identifies the percentage of total volume received in the region.

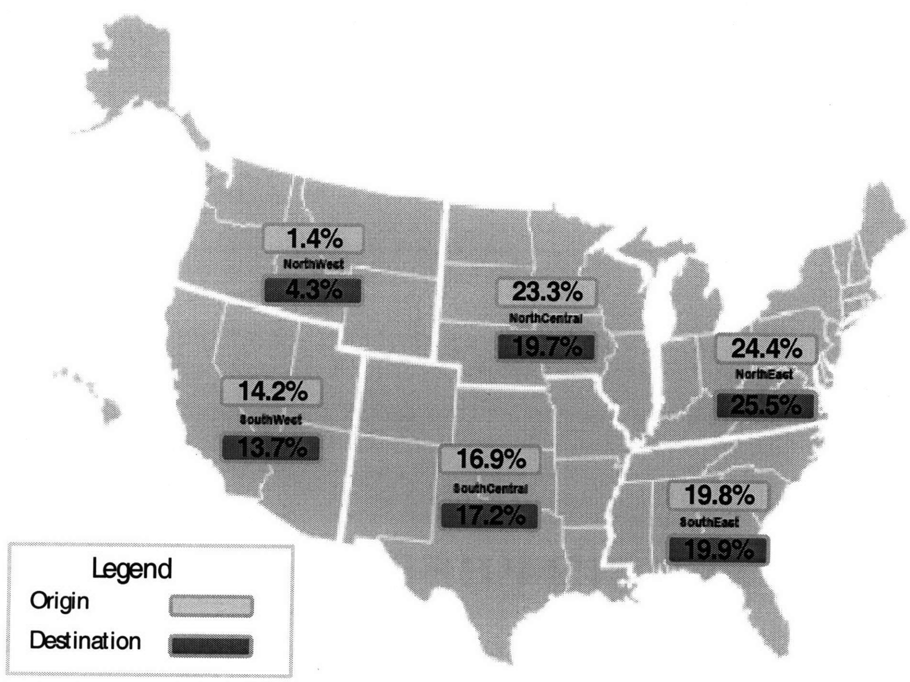

Figure 2.1: Percentage of volume by region

The map shows that the all the regions, except the North West and North Central, have a rough balance between the number of inbound and outbound loads. The North West has many more inbound loads than outbound, and the North Central has slightly less inbound than outbound. There were only eight states that shipped or received more than 4% of the total number of loads in the dataset. Even at the state level the origin and destination volume was fairly balanced with seven of the eight higher volume states having less than a 3% difference between the inbound and outbound loads. The one state that had great disparity in load balance was Ohio which shipped 14% of the dataset volume but only received 5%. Origin and destination volumes by state are listed in the appendix.

The dataset has a somewhat limited range of industries, namely paper production, manufacturing, packaging, and food and beverage. The overall food and beverage portion of the dataset represented a majority of the total loads and accounted for 72% of the volume. The other industries were fairly evenly distributed among the remaining 28% of the volume. This breakdown of volume by industry is summarized in Figure 2.2.

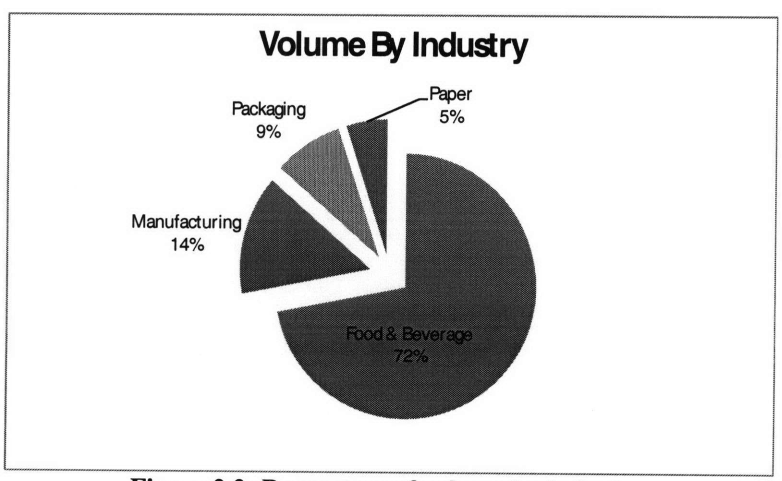

Figure 2.2: Percentage of volume by industry

The different customers had a total annual truckload transportation spend that ranged from $1.4 to over $140 million. Figure 2.3 below summarizes the amount spent on a truck load shipments greater than 250 miles. The median load cost was $898, but since the distribution is skewed heavily to the right the average load cost was $1019.

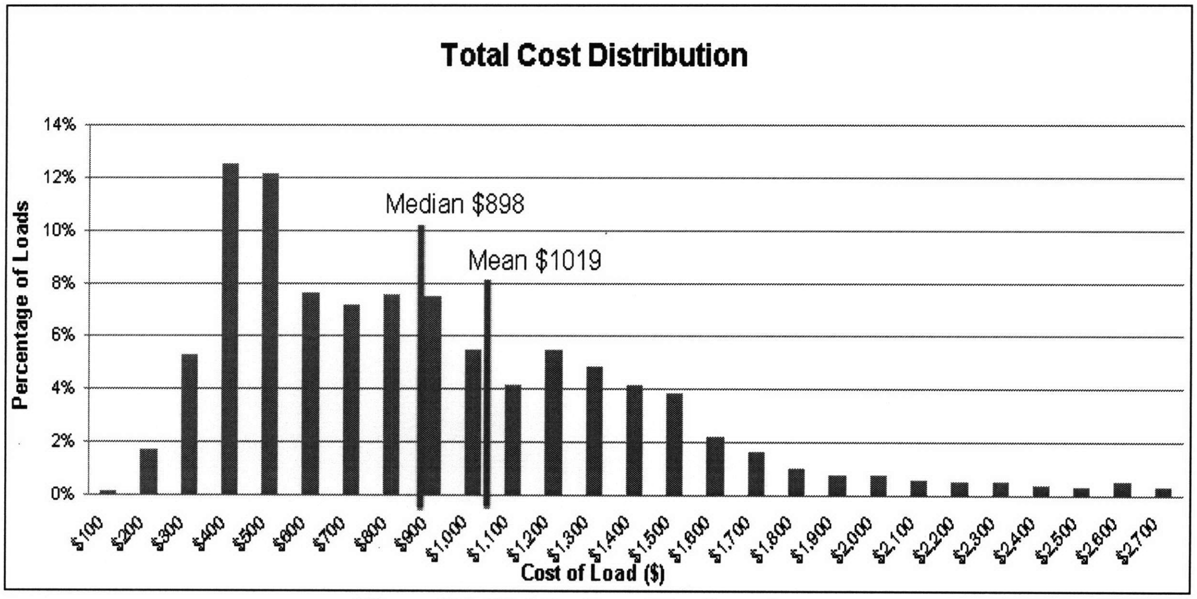

Figure 2.3: Distribution of tender rates

Looking at the customer profile, the dataset reveals that each customer has specific characteristics with regard to total number of annual loads, density of volume on specific corridors, average lead time, and routing guide depth which would lead us to expect different rates by customers for loads with similar origins and destinations. In Figure 2.4 below the largest customers in the dataset ship three to four times as much freight as the smaller customers. One would expect that the more volume a customer ships then the greater the leverage they would be able to exert with carriers to secure the most favorable rates along with more complex information technology systems to both manage their transportation network and track costs.

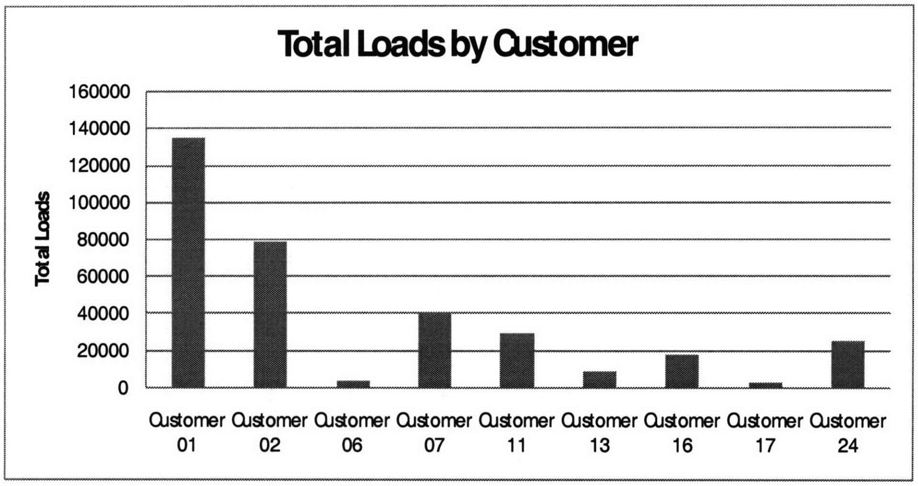

Figure 2.4: Total loads by customer

Another factor relating to customer volume is the average corridor (3-digit zip code to 3-digit zip code) volume which reflects the density of loads in a customer's transportation network. The concept behind this is that the frequency of shipping on an established route might provide carriers with additional consistency that reduces tender rejections more than just reviewing a customer's overall volume. The graph below shows some of these differences between total volume as seen in the earlier graph and the shipping density per lane. In particular customers 11 and 13 only accounted for a small minority of total volume but achieve a corridor volume density well above the average. Like total volume, we would expect corridor density to provide a cost advantage when negotiating rates with carriers and will use the direct model to help quantify corridor volume impact on total costs.

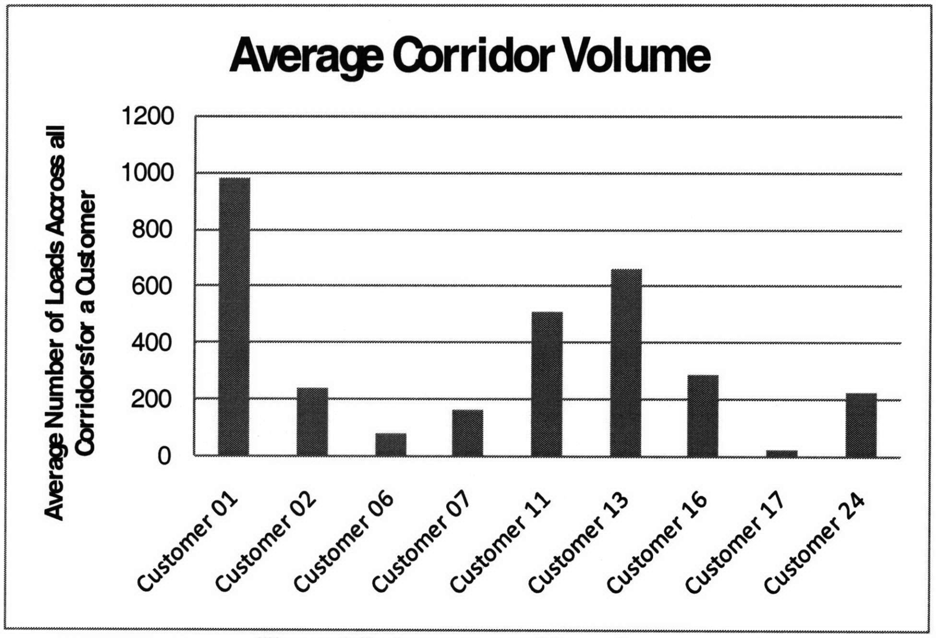

Figure 2.5: Average corridor volume

The next factor that varies widely by customer is their average tender lead time for loads. The graph below displays the differences between customers with regard to lead time (Figure 2.6). The difference in average lead time by customer is over four days with customer 13 having less than a day of lead time and customer 2 providing carriers over 5 days of lead time on average. The other customers were less extreme and only varied by a day from the average of the dataset.

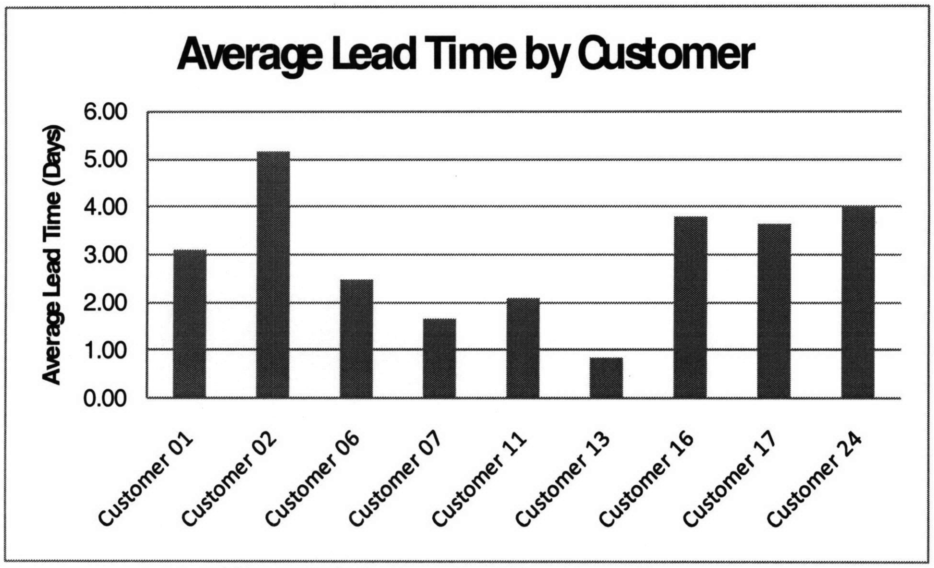

Figure 2.6: Average lead time by customer

The final difference we see in customers is the cumulative effect of the variables discussed in this section which are reflected in the average routing guide depth per customer. These differences are displayed in the graph below (Figure 2.7). From our earlier understanding of the relationship between routing guide depth and rates we would expect that customers which have to tender more often for a load on a specific corridor will end up paying more than customers that tender less on average. The challenge in this analysis is that routing guide depth will not necessarily account for the impact of volume or lane density which might also help secure better pricing. This is because the greater volume or density could help the customer negotiate better rates with carriers regardless of how often their loads are rejected.

Customers with greater volume and corridor density are also less likely to be rejected by carriers since their business is considered more important than the lower volume and lower corridor density of smaller customers. Overall it is extremely difficult to distinguish between these two effects of volume and corridor density by looking at routing guide depth or even rate per mile unless you were able to compare two customers that had different volume and lane densities but mirrored each other on the other transportation factors. For the extend of this research we have not explored those concepts more in depth and focused on quantifying the overall impact of the total volume and lane density on the total cost of an average load by taking into account all variables. With that said, it is quite possible for customer 13 to have a deeper routing guide depth but to pay less than a customer with a better routing guide depth but worse rates in their routing guide.

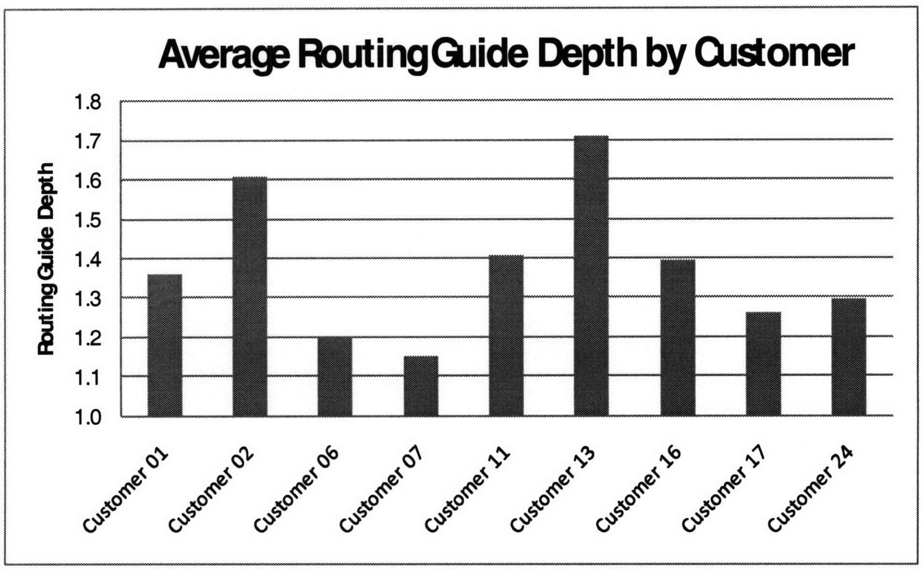

Figure 2.7: Average routing guide depth by customer

3 Routing Guide Depth

Since truckload prices are contractually fixed, cost variations are a function of how deep into the routing guide a shipper needs to go before a tender is accepted. In this chapter we report the results of our routing guide depth profile and describe our models for predicting the number of tenders per load.

3.1 Profile of Routing Guide Depth in Dataset

Routing guide depth plays a role in explaining how often a load is rejected by carriers. Each time a carrier rejects a load the routing guide depth of the load is increase by one. Examining routing guide depth in the dataset provided clues about how different variables affect acceptance rates. Lead time plays a key part on average routing guide depth. Loads that had a shorter lead time also had a much higher level of tender rejections and shippers were forced to resort to carriers deeper in their routing guides at a greater cost – this relationship had a correlation of 0.90.

The graph below (Figure 3.1) shows the cumulative percentage of volume within the dataset by routing guide depth. 78% of the volume is accepted by the first carrier. The second carrier accepts 13% of the loads which represents a large portion of the 22% of the loads the first carrier rejected. The third carrier accepts the next greatest volume of loads with 4% of the total. After the third carrier the total volume of loads accepted drops dramatically to less than 2% for the fourth carrier and then below 1% for all the carriers after that. When a carrier rejects a load there is a higher probability that the next carrier in the routing guide will also reject the load. This trend can be seen in Figure 3.2 where the probability of a carrier accepting a load drops with routing guide depth.

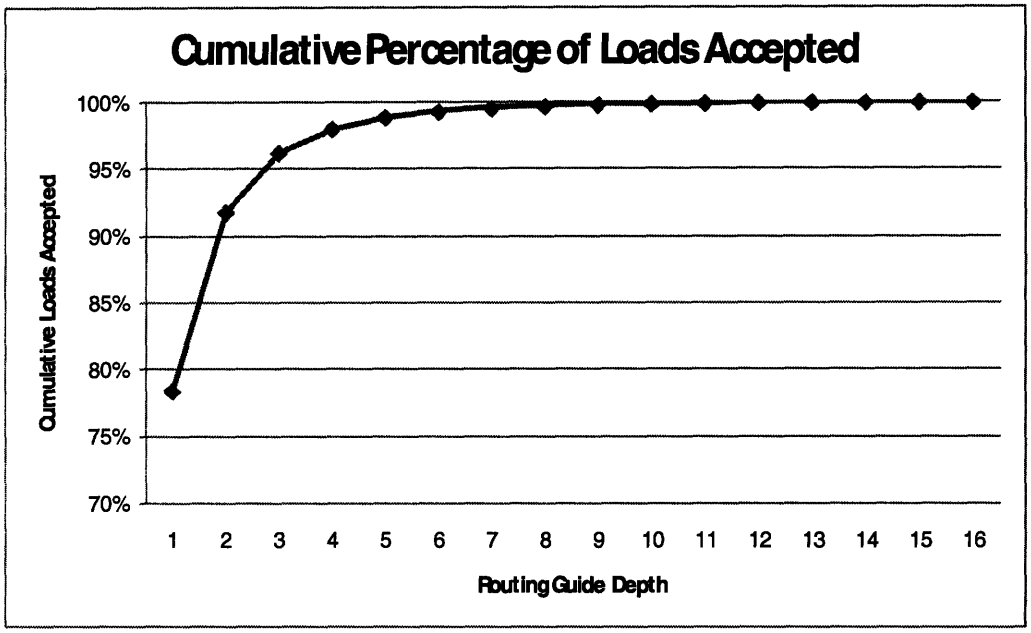

Figure 3.1: Cumulative percentage of loads accepted by tender routing guide depth

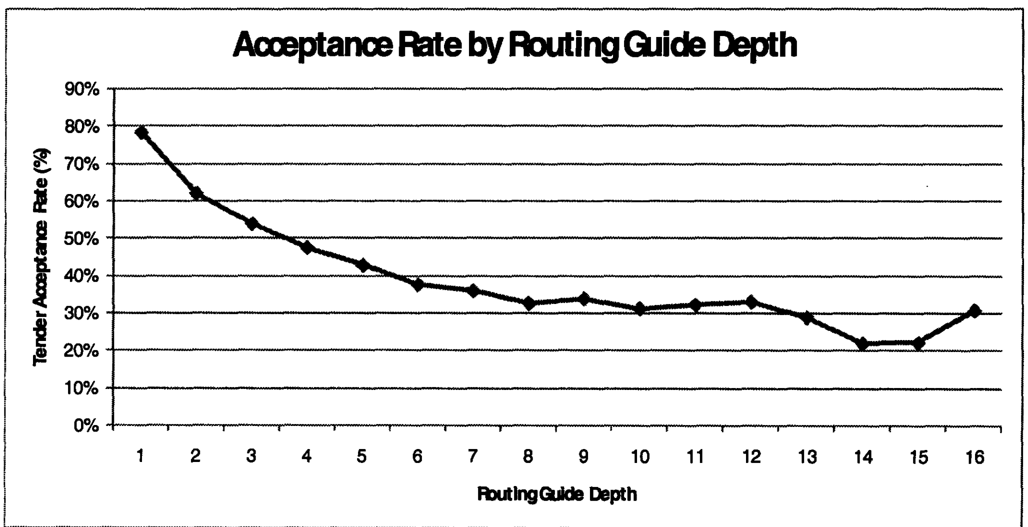

Figure 3.2: Acceptance rate by routing guide depth

For loads requiring multiple tenders, we estimated the period of time between tenders, as shown in Table 3.1. The middle column is the average time since the initial tender at each routing guide position, and the column on the right is the inferred average time since the previous tender. Each rejection adds to the time it takes for a load to be accepted and there was a fairly broad difference in the average time between tenders by routing guide depth. This is particularly pronounced with a ten hour wait between when the first carrier receives a tender and rejects the tender and when the second carrier receives the tender. We suspect this delay is often based on the original tender being sent to the first carrier in the afternoon. Since carriers typically have a number of business hours to review a tender the carrier response is not sent until the morning of the next work day.

Routing Guide Depth	Cummulative Time From Initial Tender (hrs)	Time Between Tenders (hrs)
1	0	0
2	10	10
3	14	4
4	19	5
5	25	6
6	30	5
7	32	2
8	37	5
9	40	3
10	41	1

Table 3.1: Cumulative and average time between tenders (hours)

Figure 3.3 below depicts the average level of routing guide depth for loads based on the number of days of lead time. Once again data confirm the trend that shorter lead times result in increased tender rejections and that loads with longer lead times between five and nine days are the least likely to be rejected by carriers. The curve looks very similar to the airline yield management curve presented in the introduction where there was an increase in price in the short term and an increase beyond the intermediate term which then level off again. This same curve will be seen when looking at rates by lead time and helps demonstrate the increased cost of short lead time as well as carriers' reluctance to commit to loads too far in advance since they have to balance those choices with the possibility of securing higher priced loads shorter term.

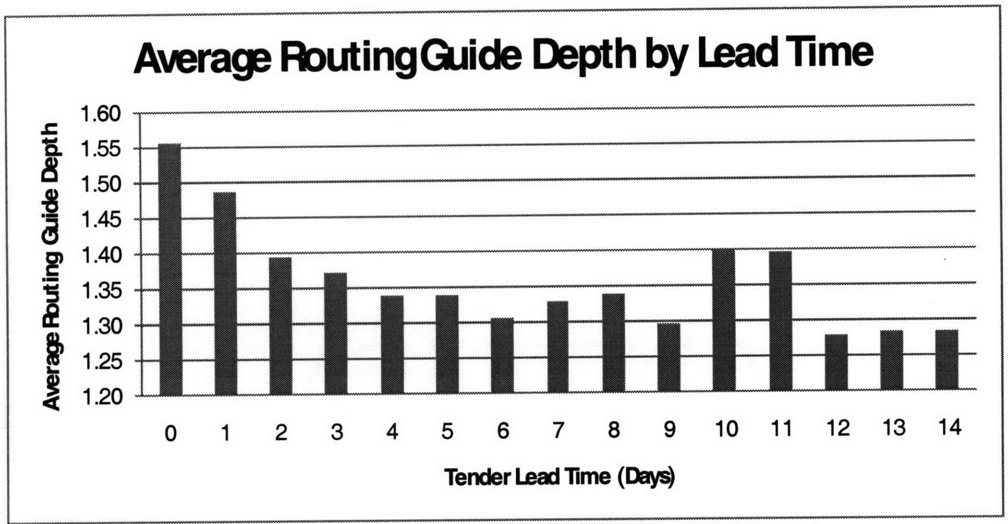

Figure 3.3: Average routing guide depth by days of lead time

The other trend that can be seen in the dataset is that for each increase in sequence there is a corresponding increase in average transportation rate per mile. Overall each increase in sequence resulted in roughly a $.06 per mile rate and this translates into an average increase of over $26 per load for each rejection in the dataset. Even though this increase in rate per mile for each step in sequence is substantial it is important to keep in mind that overall volume dropped significantly for each step in sequence.

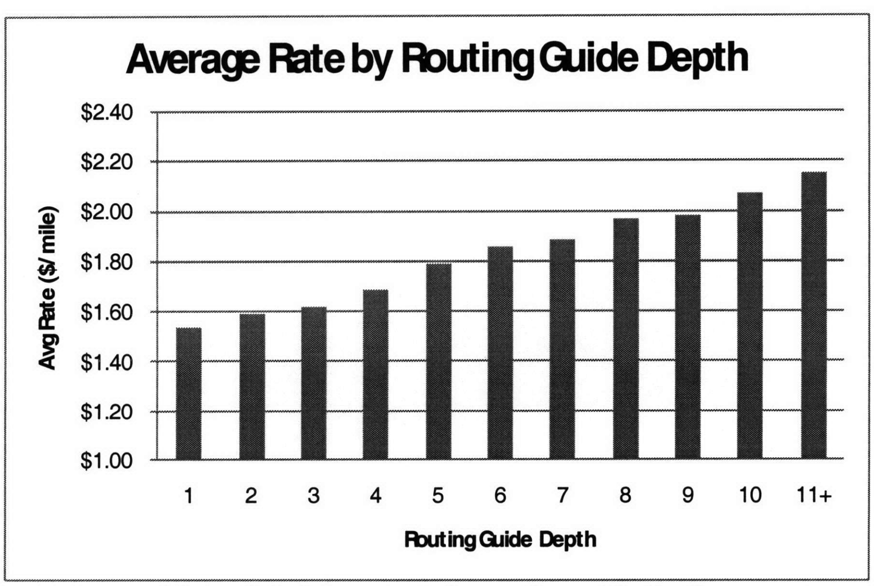

Figure 3.4: Average rate by routing guide depth

Since initial acceptance rates are so high, there are relatively few loads tendered at the highest routing guide depths. Some of the variability observed with regard to these tender routing guide depths is likely to be due to the smaller number of cases present in the data.

It becomes apparent from the dataset that the deeper a shipper has to tender into the routing guide the greater the rate per mile they will have to pay. This trend has a very steep initial slope with the most dramatic average rate increase, 7.9%, occurring when the initial tender is rejected. Each successive tender rejection equates to a 3.2% penalty in rate per mile as they solicit the next carrier to take the load. This makes intuitive sense with shippers typically placing the least cost carriers at the top of the routing guide in order to minimize total transportation spend. Figure 3.4 also emphasizes the importance shippers place on cost over qualitative service considerations assuming that the best performing carriers are not typically the cheapest as well. Unfortunately, the dataset did not contain the relative service level rankings of the loads since it would have been interesting to quantify how valuable a fractional level of service was in terms of cost; however, even from our limited data it would strongly suggest that price is the dominant factor in determining a carrier's position in the routing guide.

Looking further at expected price we used the number of times a load received a tender based on the average for the corridor on which that load travels. From Figure 3.5 below there is a trend from the dataset that the greater the lead time given to the carrier the less average number of tenders a load has on a given lane. This trend is pronounced in the 90^{th} percentile where there is a noticeable improvement in acceptance by carriers when given more than 6 days of lead time. The impact to the better percentiles was much less pronounced and the majority of acceptance improvement came from the better acceptance rate in the worst percentage of loads – the 90th percentile or higher.

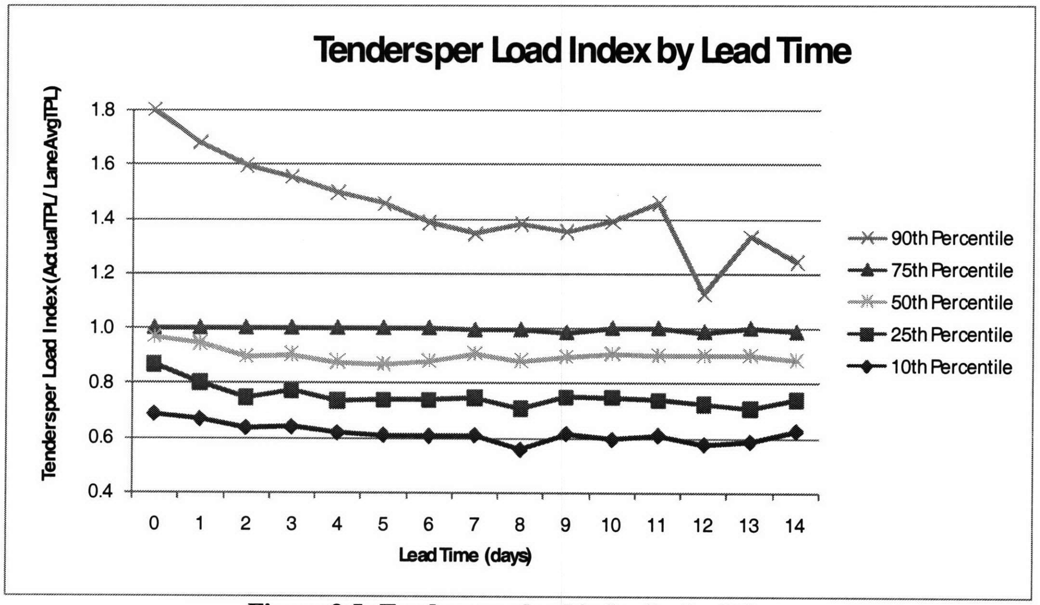

Figure 3.5: Tenders per load index by lead time

As expected from the trend presented in the last graph we can see how the more often a load has to be tendered before being accepted then the greater the impact on the rate paid per mile. The graph below shows how the cost per mile changes with respect to lead time by taking the actual cost per mile for a load and dividing it for the average cost per mile on the corridor it travels. In all percentiles of loads the performance gravitated toward the lane average as lead time increased. This is interesting since the graph above showed how the loads that required the most tenders compared to the corridor average improved substantially while the best loads seemed to be little changed with lead time. The graph below shows that in the case of the top 10th percentile and the top 25th percentile the loads actually become slightly more expensive with lead time. Overall this decrease or level performance (an increase from 0.90 to 0.94 index rate) of the top percentile of loads is more than offset by the tremendous improvement in cost per mile index for the worst percentile (1.13 to 1.05 index rate). It will be interesting to note how sensitive this data is in the direct model approach and whether lead times of 0 day prove to be cheaper than lead times of 1 day. Overall this trend of the top percentiles was probably somewhat masked in the earlier graphs showing rate per lead time since it was aggregating the entire dataset and the average rate would be heavily influenced by the worst performing loads on a given lane.

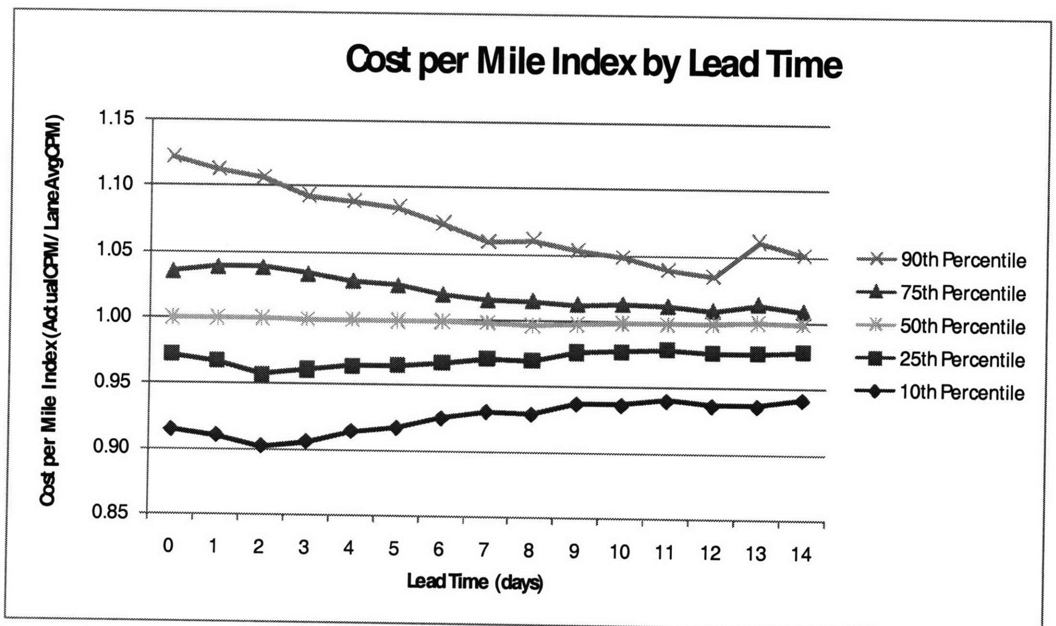

Figure 3.6: Cost per mile index by lead time

3.2 Predicting Tenders per Load

The data profiling stage revealed a 0.88 correlation between the tender sequence number of the first 15 tenders and the change in rate from the initial tender (Figure 3.8). Given this correlation, it may be possible to predict tender rates by predicting how many tenders will be required before a load is accepted and then using this number to infer an expected rate. We built several such models in order to better understand the impact of lead time on tender sequence.

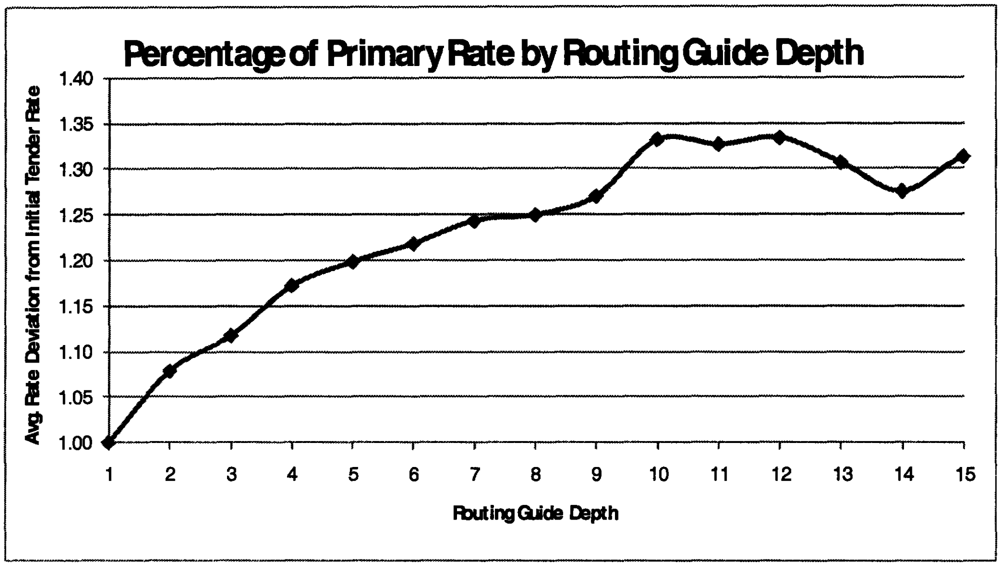

Figure 3.7: Percentage of primary rate by routing guide depth

We would like to mention two factors that may influence the performance of these models. First, since sequence number and rate do not appear to be correlated after the first ten tenders, it will be difficult to infer a tender rate even with a perfectly accurate prediction of a sequence number in this range. Second, the relationship between tender sequence and rate changes noticeably when examined on the customer level, as shown below.

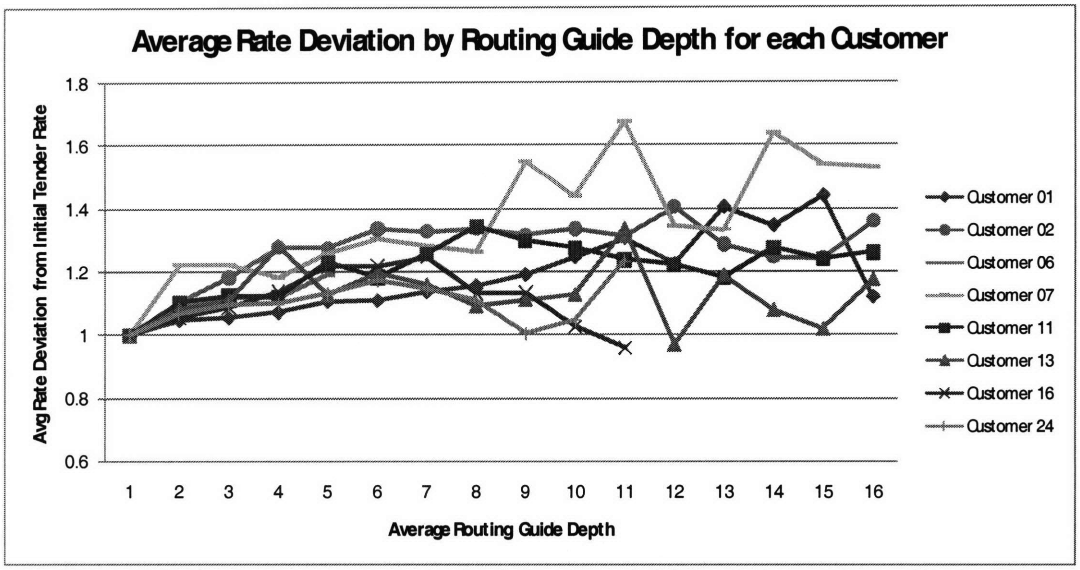

Figure 3.8: Average customer rate deviation from initial rate by routing guide depth

Since the relationship between sequence number and costs vary by customer, a model either needs to be tailored to the customer or compensate for the customer-specific behavior, routing guide characteristics, or other factors causing the variation in rates.

As shown in Table 3.2, there appears to be a nonlinear relationship between tender lead time and tender sequence number. Therefore, an integer lead time variable is likely to be ineffective at helping predict tender sequence numbers in linear models.

We built and tested several models to predict tender sequence using a common dataset containing 2921 accepted loads for a single customer over a two-week period in 2007. Using a two-week time window eliminated the effects of seasonality and economic cycles in the trucking industry, while still providing enough data to determine if the approach was feasible. Loads from 16 origin states and 44 destination states were represented in the chosen dataset. The sequence numbers of the loads and their probability in the dataset were as follows:

Routing Guide Depth	Number of Loads	Prior Probability
1	1931	0.6611
2	524	0.1794
3	193	0.0661
4	139	0.0476
5	56	0.0192
6	23	0.0079
7	17	0.0058
8	20	0.0068
9	11	0.0038
10	6	0.0021
11	0	0.0000
12	0	0.0000
13	1	0.0003
Total	2921	1.0000

Table 3.2: Distribution of loads by routing guide depth for indirect model training data

Given our objectives of predicting tender sequence and understanding the financial impact, some of the common measurements for evaluating models are misleading. For example, error rates used in measuring classification models fail to reflect the severity of an error. For a load that was accepted at sequence 1, incorrectly predicting a sequence 2 acceptance increases the error rate by the same factor as predicting a sequence 11 acceptance. Lift charts, which are often used in measuring the effectiveness of prediction models like those used in targeted marketing, are also problematic. They highlight the usefulness of a model's highest confidence predictions, whereas we are just as interested in the effect of the model's lowest confidence predictions. Ultimately, the best evaluation method needs to involve using the predicted sequence numbers for the records in the test set and a reconstructed routing guide for the corresponding lane to calculate a predicted tender rate. These predicted rates can then be compared to the actual rates using standard statistical methods.

To aid in building and testing models, the dataset was partitioned into training, validation, and test sets as described in the methods section. The result of each modeling technique we employed is described below, grouped into prediction and classification type models.

3.2.1 Multiple Linear Regression

Each of the prediction models outputs a real number prediction of the tender sequence. Although the ordinary least squares multiple linear regression we constructed was only able to explain a third of the sequence number variability, it did suggest which variables may have predictive potential, including several of the origin and destination state variables, *LeadTimeDays*, *MeanWeeklyVolume*, and *LaneWeekCount*. The R^2 value was only 0.32, so the magnitude of the coefficients cannot be reliably accepted, however the sign of the coefficients implied that tenders from or to certain states can either raise or lower the predicted sequence number.

Input variable	Coefficient	Std. Error	P Value
Constant term	1.0670	0.1022	0.0000
LeadTimeDays	-0.1144	0.0131	0.0000
ToCA	-0.3487	0.1329	0.0089
ToFL	0.5121	0.1213	0.0000
ToGA	1.0884	0.1309	0.0000
ToOK	-0.8854	0.2904	0.0024
FromFL	1.0051	0.0985	0.0000
FromTN	0.7052	0.1110	0.0000
TenderSat	5.0792	1.1747	0.0000
MeanWeeklyVolume	0.0408	0.0094	0.0000
LaneWeekCount	-0.0102	0.0028	0.0003

Table 3.3: Multiple linear regression coefficients (R^2=0.32)

Total Sum of Squared Errors	RMS Error	Average Error
842.73	1.2002	-0.0754

Table 3.4: Multiple linear regression test dataset scoring summary

3.2.2 *k*-Nearest Neighbors

The *k*-nearest neighbors classification model appeared to be the best predictor of tender sequence number. This technique uses the concept of proximity to classify each record. The distance between each of the record's inputs and those of the other observations are calculated and the results are combined to determine the *k* most similar records. This is done for values of *k* starting at 1 to some number. The error percentage, or the percentage of records classified incorrectly, for each of the *k* values is calculated and the *k* value with the smallest error percentage for the validation dataset is selected as the optimum value of *k* for the model. For our data the error percentages for *k* values between one and ten are shown below. The optimal value of *k* is two.

k	% Error Training	% Error Validation	Optimal *k* Selection
1	1.30	28.65	
2	16.16	27.74	← Best *k*
3	19.45	29.45	
4	23.63	31.28	
5	24.18	30.37	
6	25.21	30.71	
7	25.68	30.48	
8	26.99	30.94	
9	27.60	29.68	
10	28.22	30.59	

Table 3.5: *k*-Nearest neighbors error rates by k values

Using a *k* value of 2, the model produces the results below when classifying the test dataset. Note that the model never predicts that a load will require 11 or 12 tenders because these values were never observed in the training dataset.

Class	# Cases	# Errors	% Error
1	389	29	7.46
2	117	75	64.10
3	34	25	73.53
4	16	8	50.00
5	12	8	66.67
6	7	7	100.00
7	5	5	100.00
8	1	0	0.00
9	2	0	0.00
10	1	1	100.00
13	1	1	100.00
Overall	585	159	27.18

Table 3.6: *k*-Nearest neighbors test data error summary

The confusion matrix below provides additional details by specifying the actual class of all predictions made by the model. The correct predictions are in bold. The first column of data shows the actual classes of all loads that were predicted to require exactly one tender. The table shows that 360 of these loads actually belonged to class 1, 68 belonged to class 2, 16 belonged to class 3, and so on.

	Predicted Class										
Actual Class	1	2	3	4	5	6	7	8	9	10	13
1	**360**	17	9	2	1	0	0	0	0	0	0
2	68	**42**	4	3	0	0	0	0	0	0	0
3	16	5	**9**	2	0	0	0	2	0	0	0
4	6	2	0	**8**	0	0	0	0	0	0	0
5	4	2	0	2	**4**	0	0	0	0	0	0
6	0	1	0	0	6	**0**	0	0	0	0	0
7	1	3	0	0	0	0	**0**	1	0	0	0
8	0	0	0	0	0	0	0	**1**	0	0	0
9	0	0	0	0	0	0	0	0	**2**	0	0
10	0	1	0	0	0	0	0	0	0	**0**	0
13	0	1	0	0	0	0	0	0	0	0	**0**

Table 3.7: *k*-Nearest neighbors test data confusion matrix (correct predictions in bold)

The sequence-1 error rate of 7.46% is higher than many of the other models constructed, but the subsequent error rates are lower. Therefore, we suspect this model's prediction of the total cost of a group of tenders will be the most accurate of all the models.

3.2.3 Additional Models

We tried building models using several additional techniques, including classification tree, naïve Bayes, neural network, and discriminant analysis. None of the models produced using these techniques produced promising results.

When scoring the test data, the classification tree we constructed never predicted a class higher than sequence 4. This suggests that a larger sample or a more evenly weighted sample would improve the model. The top nodes of the pruned classification tree and error reports are shown in the Appendix.

The naïve Bayes model we built had the highest error rate of all the models. This may be because the derived probabilities were inaccurate due to the small amount of training data it was provided.

The neural network classification model did a poor job of classifying even the training data with the exception of sequence 1 (see Appendix for complete test results). Since sequence 1 has the highest probability this could be achieved even with a naïve model. This suggests that a better way to train the models would be a weighted sample in which the distribution of sequence number is more balanced. We also noticed that the model even performed poorly on the training set, instead of overfitting the data as is often the case. This suggests that the available inputs were not capable of explaining the variety of tender sequence outcomes. Additional inputs may be required to improve the performance of any of the models.

The discriminant analysis model we constructed had a 35% test set error rate for tender sequence 1 (see Appendix for complete test results). This was disappointing since this technique is known to do well on small datasets. The technique however relies on an assumption that the output is approximately normally distributed. Since the tender sequence data is heavily skewed, this assumption is obviously violated. It is possible that the results could be improved by using the logarithm of the sequence number to reduce its skewness.

3.2.4 Summary of Models to Predict Tenders per Load

Although the *k*-nearest neighbors model appeared to be the most effective, none of the models were capable of predicting the number of tenders per load correctly for more than 75% of the tenders in the test dataset. However, building and evaluating the models taught us several interesting lessons. Decision rules in the classification and regression trees and the signs of the linear regression helped us think more about the contribution of specific inputs. We also recognized the need to recode continuous variables like *LeadTime* so that they can be better used in linear models. The results also revealed several areas in which additional work could be done. We think many of the models could be improved with a larger training set. In future work, we would also like to try weighting the training set to make higher sequence numbers more common.

4 Lead Time Modeling

With an understanding of how routing guide depth affects cost, this section uses an approach of dataset profiling and linear regression to better quantify the impact of lead time and other factors on the cost of a load.

4.1 Profile of Lead Time in Dataset

By using the profile of lead time in the dataset it was possible to gain a better understanding of expected lead time impact and trends. Using the graph below one can see that two days of lead time is the most common at 20% of the loads but was closely followed by one day of lead time at 19% (Figure 4.1). This trend continued in a linear manner until overall volume by day of lead time fell below 1% at eight days of lead time. Interestingly the number of loads with zero days of lead time accounted for only roughly 6% of the total volume of loads. The average lead time of the dataset was 3.4 days with a median of 2.8 days.

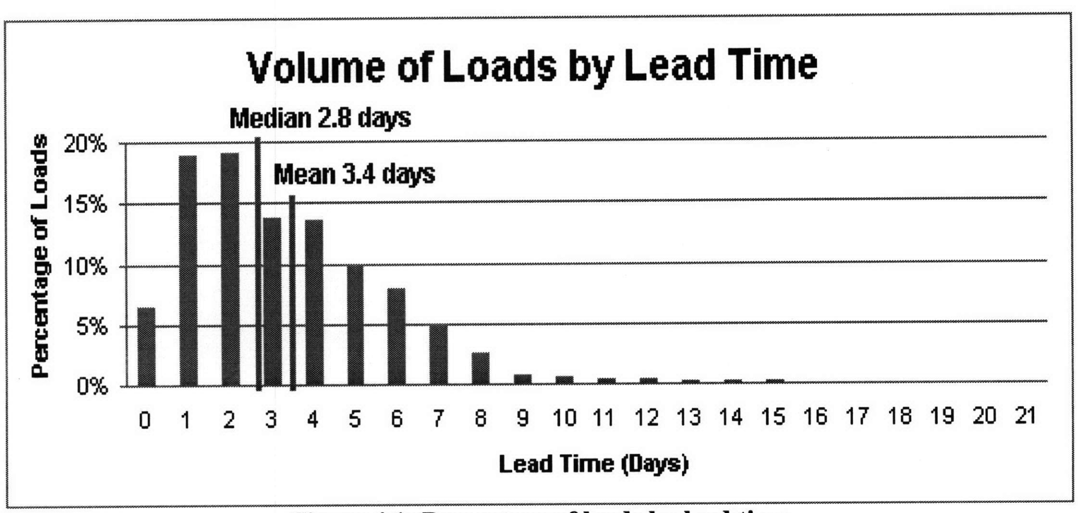

Figure 4.1: Percentage of loads by lead time

The average rates were fairly flat for the first few days of lead time excluding loads with zero days of lead time. The average rate then dramatically drops for lead times beyond five days. Even with the initially flat change in rates at the left of the graph, there is still enough volume beyond four days of lead time to make the $.06 reduction in rates for each additional day at the longer lead times significant. This translated into a savings of $41 for each additional day of lead time between five and eleven days of lead time without considering any other variables. This trucking rate curve loosely resembles the expected pricing curve for airlines as discusses in the literature review and makes sense that carriers accept a higher percentage of loads a week in advance of pick up since much of their trucking capacity is uncommitted at that point. This is also the exact same pattern observed when looking at routing guide depth by lead time. Trying to tender inside of a week runs the risk that the carrier has already committed their assets and will therefore reject tenders at a higher frequency than at the longer lead times. Finally, we see an increase in rate as we approach two weeks lead time where carriers are hesitant to plan loads that far in advance. This hesitation could be due to inability to accurately forecast location of assets that far in advance or reluctance to commit to the lower prices in advance when rates might be more favorable for them in the shorter term. The more robust model in the next section will explore the sensitivity of rates to lead time at the greater extremes by isolating other variables could be having an impact in order to obtain a truer understanding.

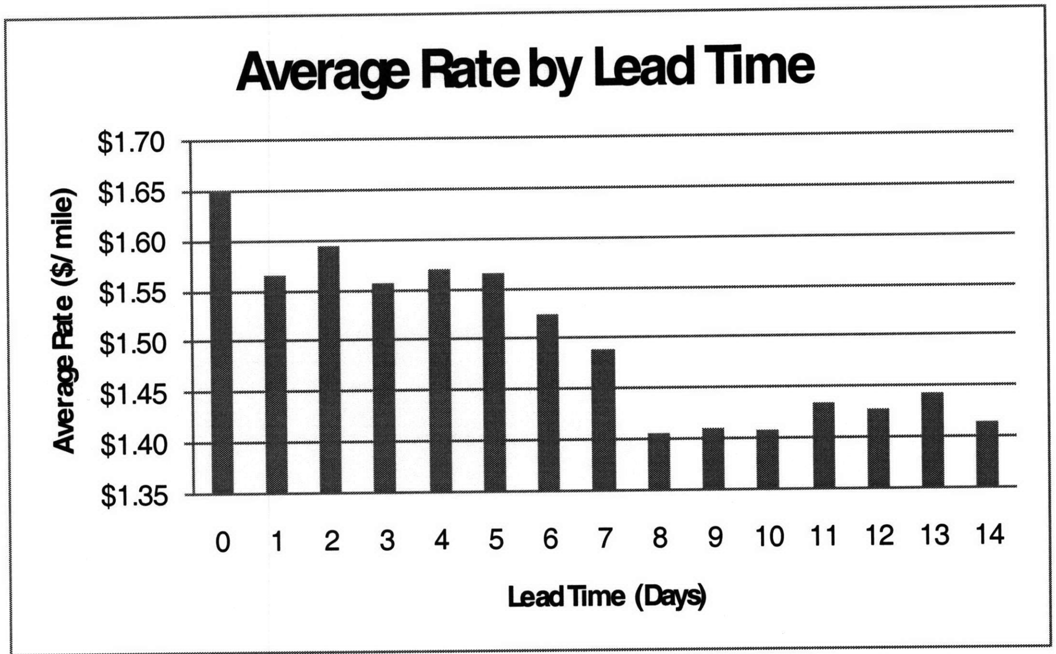

Figure 4.2: Average rate by lead time

Part of our analysis was discerning how sensitive load costs were to different groupings of lead time in the regression models, which was an iterative process. We started with lead time in smaller hourly groupings and then expanded to larger groupings as they proved statistically significant and meaningful.

By doing this detailed analysis we found that lead time had a more general effect on rates and the more granular levels of grouping only provided marginally more detail at the expense of making the model substantially more complicated. For instance, for longer lead time values the change in costs was small enough that a grouping of all loads tendered beyond 12 days of lead time accurately captured the trend. Even though there were some increases and decreases in the hourly groupings on the shorter end of the lead time spectrum, these also proved to be fairly linear with a flat slope between 17 and 59 hours. The 17 to 59 hour lead time group was used as the base case in the transportation model and loads tendered before that time period incurred a cost penalty and loads tendered after that time period received a cost discount.

Using the transportation model confirmed what we deduced from the dataset profiling that the savings is fairly flat around the 17 to 59 hour window but then increases dramatically after 5 days until leveling out at 12 or more days of lead time. The model supported the assumption that there was enough volume of loads at even the greater lead times to make those reduced rates statistically meaningful. The next step in our analysis was to use the insights we found on the impact of lead time and expand the transportation model to account for all relevant factors that affect cost and then determine how lead time plays a role in the broader model.

4.2 Impact of Lead Time in the Transportation Model

We used an ordinary least squares linear regression model to explain the impact of independent variables on cost per load. The variables in the model are summarized in the table below and each variable is discussed in detail in the subsequent sections.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Constant	$ per load	$ 149.90	0.000	$145.9531	$153.8529
Distance	$ per mile	$ 1.17	0.000	$1.1674	$1.1715
Origin	Origin State ($ per load)	Various			
Destination	Destination State ($ per load)	Various			
Corrider Volume (per year)	Daily >150 loads ($ per load)	Base Case			
	Weekly = 52-149 loads ($ per load)	$ 1.69	0.063	($0.0884)	$3.4636
	Monthly = 24-51 loads ($ per load)	$ 21.27	0.000	$18.7741	$23.7734
	Annual <24 loads ($ per load)	$ 61.51	0.000	$59.1380	$63.8824
Carrier Size	Small/Medium	Base Case			
	Large	$ 22.12	0.000	$20.6008	$23.6403
	Extra Large	$ 77.91	0.000	$74.3564	$81.4545
Tender Day	Sunday	N/A			
	Monday	$ (0.38)	0.000	($2.4503)	$1.6918
	Tuesday	Base Case			
	Wednesday	$ (2.07)	0.045	($4.0934)	($0.0430)
	Thursday	$ 4.05	0.000	$1.9939	$6.1001
	Friday	$ 5.94	0.000	$3.7767	$8.0985
	Saturday	$ 49.23	0.000	$29.0772	$69.3746
Pick Up Day	Sunday	$ 16.55	0.000	$11.4502	$21.6450
	Monday	$ 0.61	0.586	($1.5940)	$2.8191
	Tuesday	Base Case			
	Wednesday	$ (1.00)	0.366	($3.1647)	$1.1666
	Thursday	$ 2.29	0.050	($0.0038)	$4.5923
	Friday	$ (1.54)	0.153	($3.6482)	$0.5721
	Saturday	$ 22.98	0.000	$19.6140	$26.3476
Lead Time	0 to 8 hours	$ 24.26	0.009	$20.7965	$27.7145
	9 to 16 hours	$ 17.48	0.007	$14.2117	$20.7408
	17 to 59 hours	Base Case			
	3 to 5 days	$ (15.34)	0.000	($17.1003)	($13.5713)
	6 to 8 days	$ (24.51)	0.000	($26.6863)	($22.3302)
	9 to 11 days	$ (38.68)	0.000	($43.5199)	($33.8357)
	12 or more days	$ (47.10)	0.000	($53.1669)	($41.0303)
Adjusted R^2 Goodness of Fit		0.905			

Table 4.1: Transportation model results

Overall the model provided insight into how important the different factors are that influence transportation costs. Table 4.2 below shows how each input variable contributed to the model's predictions. We calculated these values by multiplying the number of loads that met the specific criteria by the absolute value of the cost impact for the criteria. For example, for the large carriers we calculated the number of loads that had a large carrier and multiplied that by the absolute value of the cost impact (62,665 loads * Absolute Value ($22.12) = $1.39M). This same process was repeated for each variable and the results were divided by the total of all variables.

Overall the transportation model is almost entirely influenced by the baseline variables – miles, origin, and destination. All the other variables besides the baseline variables, to include lead time, only make up roughly 1% of the total influence on transportation spending.

Input Category	% Contribution
Miles	65.02%
Destination	12.51%
Constant	11.71%
Origin	9.90%
Pickup Day of Week	0.19%
Tender Day of Week	0.19%
Lead Time	0.18%
Carrier Size	0.17%
Frequency	0.14%
Total	100.00%

Table 4.2: Impact to cost by variable in model

Not surprisingly, the average impact tends to minimize the greater impacts seen at the extremes. This is similar to the number of tenders per load evaluated by percentile where the best loads had very little change with routing guide depth or lead time but the worst loads improved substantially. That scenario helped to highlight the importance of focusing on preventing the extreme negative outliers. In looking at the different controllable factors that only made up roughly 1% of the average load cost we were able to construct a best and worst case scenario combining the extremes of each variable to understand the overall impact on the cost of a load. Table 4.3 below summarizes the differences by listing two values of each variable that correspond to the most extreme cost bonus or penalty. The average 1% impact of the variables proved to have a much greater financial impact. The best case scenario provided $50.71 in credit on average cost per load while the worst case scenario resulted in a $235.89 cost penalty. These two extremes combined to make a $287 difference on a load that costs an average of $1019. The transportation model discussion found in the following chapters will help quantify these impacts for each of the variables across each of their possible ranges to allow better business policy planning.

	Best	Bonus		Worst	Penalty
Corridor Volume	Daily	$ -		Annually	$ 61.51
Carrier	Small/Medium	$ -		Extra Large	$ 77.91
Tender Day	Wednesday	$ (2.07)		Saturday	$ 49.23
Pick Up Day	Friday	$ (1.54)		Saturday	$ 22.98
Lead Time	12 Days or More	$ (47.10)		Less than 18 Hours	$ 24.26
Total		$ (50.71)			$ 235.89

Table 4.3: Best and worst case controllable model variable scenarios

4.2.1 Baseline Transportation Factors

The most important baseline variables that we used in evaluating the impact of lead time continue to provide the framework upon which additional variables were added and evaluated. The impact of the three most explanatory variables, distance, origin, and destination, is summarized in Table 4.4 below. The rate per mile is multiplied by the total number of miles, the origin and destination represent various penalties and bonuses added to each load based on which state the load started and ended.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Constant	$ per load	$ 149.90	0.000	$145.9531	$153.8529
Distance	$ per mile	$ 1.17	0.000	$1.1674	$1.1715
Origin	Origin State ($ per load)	Various			
Destination	Destination State ($ per load)	Various			
Adjusted R^2 Goodness of Fit		0.905			

Table 4.4: Baseline variables for transportation model

These variables combined to provide an adjusted R^2 value of 0.905, in other words they explained 90.5% of the variability of the costs in the dataset. It was interesting to note that the origin and destination variables provided different levels of accuracy in explaining the dataset. The destination had proved to a better explanatory variable with an adjusted R^2 value of 0.852 was achieved using just the destination and number of miles as inputs. The origin failed to explain as much of the relationship in truckload pricing with an adjust R^2 value of 0.807 when analyzed with total miles.

4.2.2 Understanding the Impact of Lead Time

Now that we have completed the model with respect to all significant transportation variables how does lead time affect the overall cost of the load? From the transportation model we found that lead time is indeed statistically significant and has an impact on the price customers pay to ship loads. According to the transportation model, lead time can make a substantial difference in cost per load. Loads with a short lead time (0-8 hours) add $24.26 to the predicted cost of the load while loads with longer lead time (12 or more days) reduce the predicted cost by $47.10 per load – overall a $71.36 savings by taking advantage of longer lead time.

Even making smaller changes in lead time can make a difference to customers where increasing lead time from less than two days to over three days would improve the carrier acceptance rate and save an average of $15.34 per load. This $15.34 cost reduction on its own might seem trivial when compare against the $1019 average cost of a load but if a customer can improve their average lead time for all loads this savings can help reduce their transportation spend. For instance customer 7 had a below average lead time of less than two days in the dataset which allows the potential to improve their carrier notification by a day or two. Since customer 7 ships over 40,000 loads in a year the savings of the longer lead time would be equal to $613,000 per year by changing their lead time transportation policies without any capital investment or network redesign.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Lead Time	0 to 8 hours	$ 24.26	0.009	$20.7965	$27.7145
	9 to 16 hours	$ 17.48	0.007	$14.2117	$20.7408
	17 to 59 hours	Base Case			
	3 to 5 days	$ (15.34)	0.000	($17.1003)	($13.5713)
	6 to 8 days	$ (24.51)	0.000	($26.6863)	($22.3302)
	9 to 11 days	$ (38.68)	0.000	($43.5199)	($33.8357)
	12 or more days	$ (47.10)	0.000	($53.1669)	($41.0303)
Adjusted R^2 Goodness of Fit		0.905			

Figure 4.3: Lead time impact in transportation model

Overall, the model behaved as expected with the additional variables helping to better account for the total cost of the load. The model helped better understand the impact at the extreme short and longer periods better than the dataset profiling and was better able to attribute some of the impact at those extremes to other variables and not just lead time. Although the model returned a more conservative impact of lead time at the extremes than just looking the dataset or regressing lead time independently it proved that this factor makes a difference in transportation pricing and should be considered when planning transportation policy.

4.2.3 Carriers

The next variable that our research explored was how transportation pricing was affected by carrier size. Carriers were divided into different groups using the annual Transport Topics rankings – carriers in the top ten we classified as extra large, top 11 to 50 were classified as large, carriers ranked 51 to 100 were considered medium sized, while carriers that were not large enough to be ranked in the top 100 were listed as small. In this portion of the analysis we quantified the overall average price per mile of the different sized carriers along with overall volume of loads handled by the different sizes of carriers. We also separated the difference in average routing guide depth by carrier size to better understand how deep the different sized carriers tended to be in the routing guide. This was important in order to understand if larger carriers are generally used deeper in the routing guide than primary carriers on lanes. This would make sense if larger carriers are less likely to refuse loads at a higher cost but would bias a per mile rate against the larger more reliable carriers.

By profiling the dataset we can see that the extra large and large carriers had a higher per mile rate than the smaller carriers (Figure 4.4). This is not surprising since size of carrier provides very little cost advantage at the individual truck level while incurring significant overhead expenses at the larger carriers. One would expect some of this additional overhead to be offset with increased economy of scale with the larger carriers but the smaller carriers not only greatly outnumbered the larger carriers but also handled considerably more load volume (Figure 4.5). It is possible that the larger carriers also tend to service areas in less dense parts of the country compared to smaller carriers or even provide service to more remote parts of states than local carriers and therefore charge more for this coverage. Another consideration in comparing average rate by carrier is the relative level of service provided by the carriers is exempt from this analysis. It is possible that the larger carriers provide additional services such as system integration, reporting, truck level metrics, or account management while the smaller carriers do not and this helps reflects some of the differences in prices.

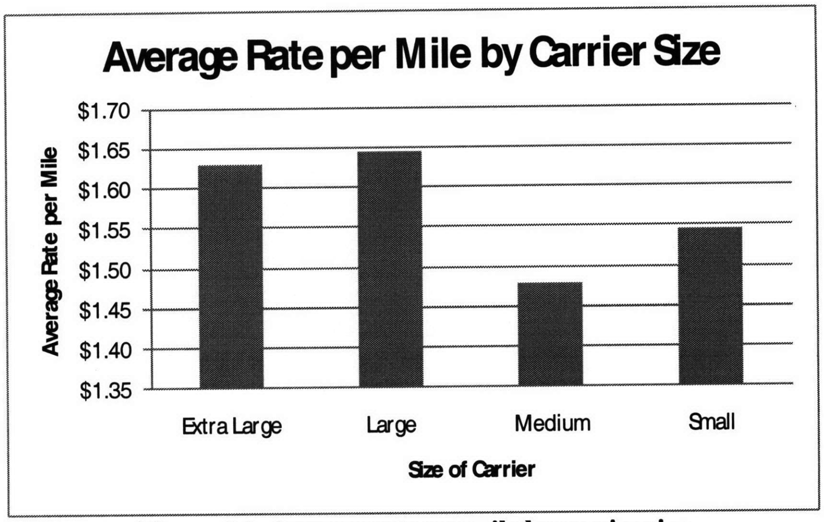

Figure 4.4: Average rate per mile by carrier size

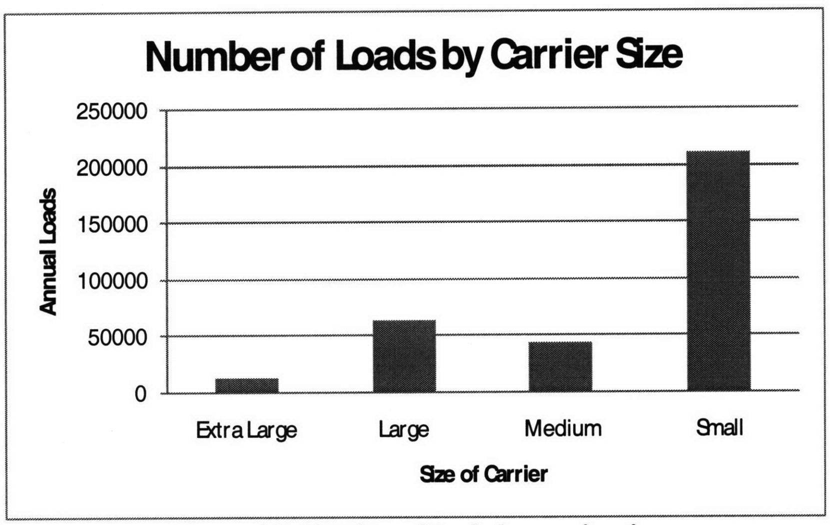

Figure 4.5: Number of loads by carrier size

The carrier size and pricing difference could be attributed to how customers treat carriers in their routing guide. The dataset profile indicates that larger carriers might have a slight disadvantage compared to the smaller carriers since their average routing guide depth was indeed 0.1 higher. Overall this appears to a fairly marginal difference and would be difficult to expect this difference to account for the roughly $.12 per mile additional cost the bigger carriers charge.

Another trend one can see from the dataset is that the longer the lead time a carrier receives on average lower the routing guide depth. This makes sense that the more time a carrier is given the more likely the carrier is to accept a tender. It is interesting to note in the graphs below that the larger carriers are given less lead time on average than the smaller carriers and there is a corresponding increase in routing guide depth. Like routing guide depth, lead time might help explain some of the differences between the different sized carriers but it is again unlikely that such a small difference in routing guide depth could account for such large difference in rate per mile.

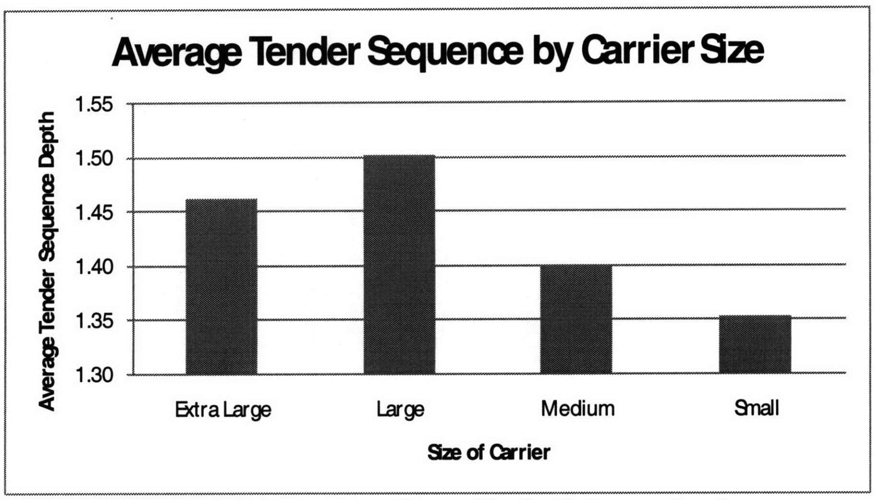

Figure 4.6: Average routing guide depth by carrier size

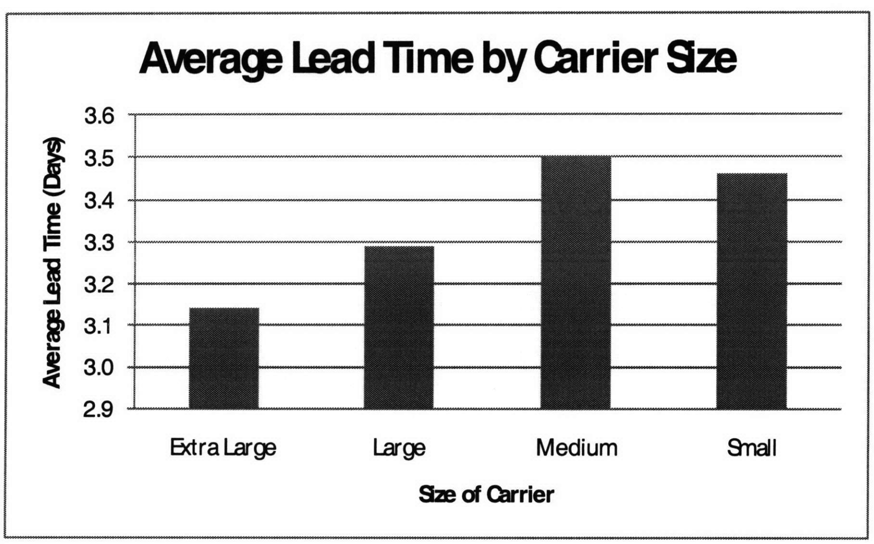

Figure 4.7: Average lead time by carrier size

Using the transportation model we find that, as suggested by the dataset profile, loads accepted by larger carriers have a higher predicted cost than those accepted by smaller carriers. Even after accounting for differences in origin and destination, lead time, corridor volume, and customer specific factors, carrier size makes a substantial difference on what a shipper would expect to pay for a load. Table 4.5 below summarizes these penalties with shippers paying an additional $71.91 per load for an extra large carrier and $22.12 additional per load for a large carrier when compared to using a small or medium carrier.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Carrier Size	Small/Medium	Base Case			
	Large	$ 22.12	0.000	$20.6008	$23.6403
	Extra Large	$ 77.91	0.000	$74.3564	$81.4545
Adjusted R^2 Goodness of Fit		0.905			

Table 4.5: Carrier size impact in transportation model

4.2.4 Sensitivity to Carrier Availability

Part of our analysis focused on how transportation rates are affected by the availability of carriers in comparison to demand and whether this effect exaggerated or mitigated by lead time. This section looks at the broad impact to the dataset during times of limited and excess trucking capacity. We used the Morgan Stanley Freight Index (Morgan Stanley 2008) to provide an industry-wide benchmark to compare periods of tight and loose carrier availability. The index compares the demand for the number of trucks in a week against the number of trucks available in that given week. Both the supply and demand values are graphed separately by Morgan Stanley and then they calculate the difference between the two to derive the freight index which shown in the graph below. Both 2004 and 2005 indicate a large demand for trucks that could not be met by the current levels of trucks on the road. Typically this would mark a period of increased transportation pricing where shippers compete more aggressively for the limited resources. It could also mark a period of increased attractiveness of alternative forms of transportation such as rail or sea freight to compensate for the limited and more expensive trucking assets.

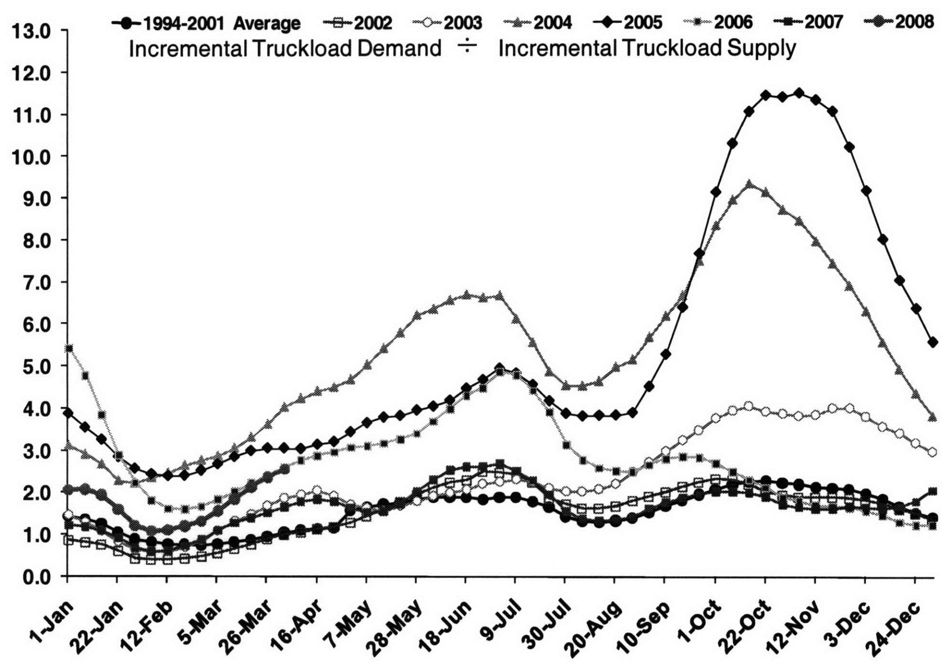

Figure 4.8: Morgan Stanley freight index (Morgan Stanley 2008)

In our dataset we looked to see how the load records were broadly impacted at various times during 2007 in comparison to the freight index. The graph below shows how a four-week moving average routing guide depth of accepted loads performed in comparison to the index. The left vertical axis displays the Morgan Stanley Freight Index values while the right vertical axis shows the average routing guide depth of loads during the weekly time frame.

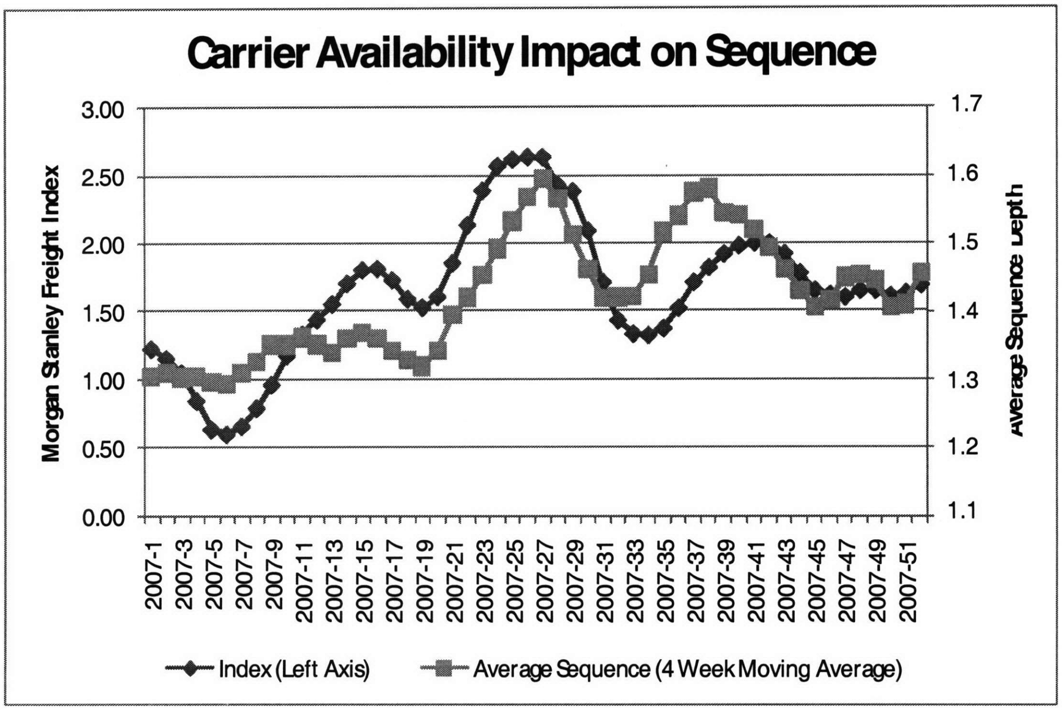

Figure 4.9: Comparison of freight index and routing guide depth

As expected, the graph depicts an increase in routing guide depth, which indicates shippers were tendering deeper in their carrier preferences, due to increased selectivity of carriers or lack of available assets by the carrier. Overall the correlation of these two variables was closely correlated at 0.73. The models in the next section will help better understand this relationship by examining each of the tender records independently to the daily freight index.

From our earlier high level discussion of the dataset we saw that the deeper a shipper is forced to tender in their routing guide there is an increase in the amount they will spend for transportation. The graph below loosely relates that expectation to what we experience in the accepted loads. The left vertical axis displays the average rate per mile paid to transport a load and the right vertical axis displays the average routing guide depth of those same loads. Like the above graph we saw that in periods of limited carrier availability the routing guide depth increased and there was an increase in transportation freight rates. Overall the correlation at the aggregate level of taking weekly average rate per mile compared to the weekly average routing guide depth had a strong correlation of 0.88 in the dataset. The graph below supports that general trend between the two variables.

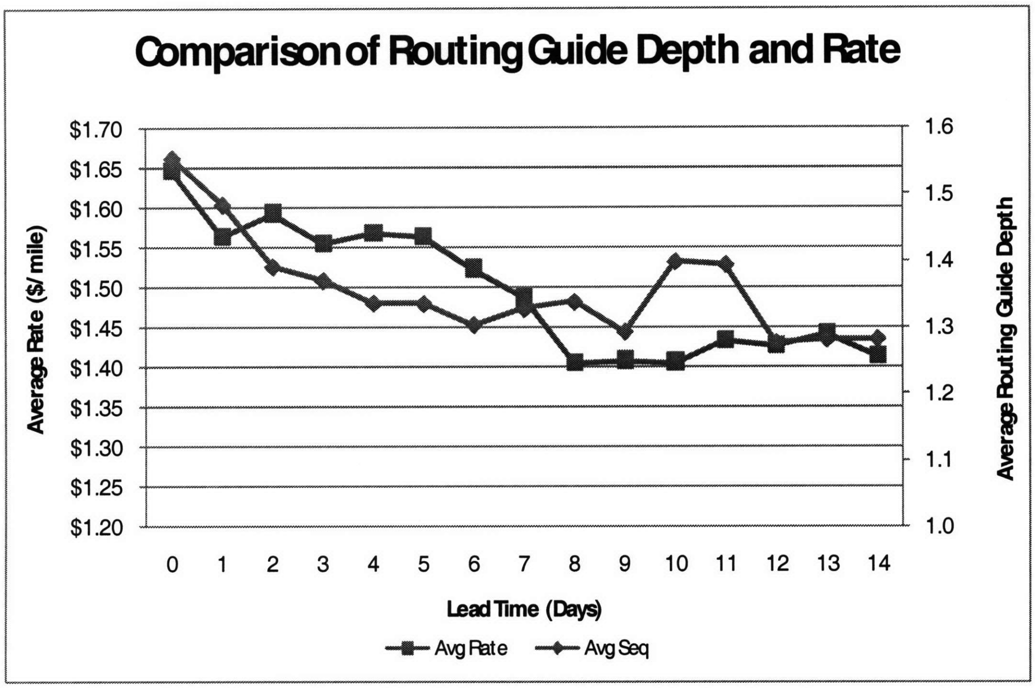

Figure 4.10: Comparison of lead time to routing guide depth and average rate

When looking at the impact of freight index in the transportation model the total cost per load actually declined when the freight index showed a contraction in asset availability. This was unexpected given the trends observed in the above graphs, but we have not yet been able to learn from Morgan Stanley if they how sensitive freight rates are to the index or what would cause a result in the opposite direction. Part of our own consideration on this outcome was that the total cost per load was very minimal at only $2.82 per load when tendered in the week where the index was greater than 1.61. Even looking at the more extreme values of the index of over 2.20 (an index value representing only 10% of loads) the cost impact was only $5.60. Both of these proved to be statistically significant in the transportation model but neither was great enough to believe it really explained an underlying process in the carrier availability.

Overall it appears that the Morgan Stanley index is fairly well correlated with routing guide depth at 0.74 and that the correlation between routing guide depth and rate per mile is correlated at 0.88 but there is less of a correlation directly between the index and the rate per mile. The other problem is that routing guide depth is fairly responsive to trends but that contractual rates might be more robust and will only be updated annually making the index less useful at explaining the impact in the short term. Finally, the year 2007 was a fairly mild year for the index compared to the years before where the more extreme index values could better explain increases or decreases in rates.

The next step in exploring the effects of the index was to subdivide the dataset by periods of extreme values for the index in 2007. We looked at the dataset with regard to the transportation variables when the index was above 2.20 and below 1.20 (these index values each represented 10% of the loads during times of tight and loose availability and the models had an R^2 value of 0.910 and 0.904 respectively) to determine the impact on cost per load. From this analysis we saw that the short lead time and carrier size became increasingly more sensitive. In particular, the penalty for the extra large carrier decreased from an average penalty of $60.46 in the full dataset to $32.17 in times of reduced carrier availability. In times of increased carrier availability the penalty for the extra large carrier increased from the average of $60.46 to $181.42. This change would make sense with smaller carriers being more sensitive to idle tractors in the times of reduced demand

and either discounting rates or increasing their acceptance rates while the larger carriers are more prone to maintain the status quo. This would explain the widening gap between the different sized carriers as large carriers become even more expensive in relation to the smaller carriers in times of reduced demand and excess capacity. The cost impact would also support the concept that in periods of increased demand the smaller carriers become more selective and refuse tenders at a much higher frequency. This would make the larger carriers seem more attractively priced in comparison during these times.

The impact to lead time was also substantial when looking at the extremes in the dataset with regard to the freight index. The average impact of short lead time (0 to 8 hours) using the full dataset in the transportation model was $34.16 per load. In times of increased demand the penalty for short lead time increased to $46.19. Likewise, in times of decreased demand the penalty for short lead time decreased to $17.39. These results support the importance of carrier availability when the freight index is at the extremes. It was unfortunate that incorporating the carrier availability variable directly in the model, rather than subdividing the dataset and rerunning the model, did not provide a greater understanding of the impact. This seems to be an area worthy of future research on how to better incorporate the impacts of the freight index into a transportation pricing model along with understanding how useful the index is at predicting macroeconomic trends and transportation company equity valuations.

4.2.5 Day of Week

Another factor we looked at in the dataset was when loads were tendered and when they were picked up. The tender day volume proved to be very consistent across days of the week with an almost complete absence of tender activity on the weekends.

This makes sense since most companies plan loads and tender with carriers during the normal work days and would most likely use the spot market for emergency shipments without lead time on the weekend. The average rates by tender day also reflect the consistency seen in scheduling loads across the days of the week and rates tended to be fairly flat with Monday being the most expensive and Friday being less expensive – overall magnitude of difference between the two days was roughly $.04 per mile.

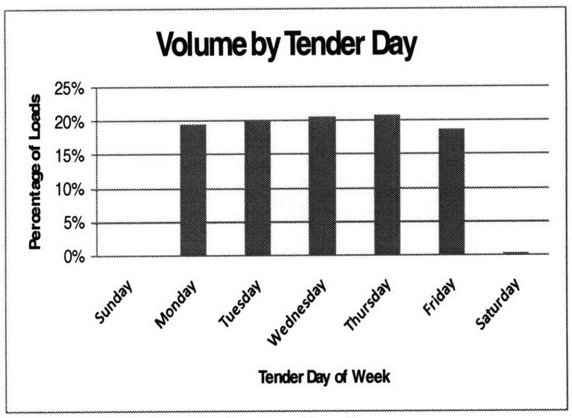

Figure 4.11: Loads by tender day

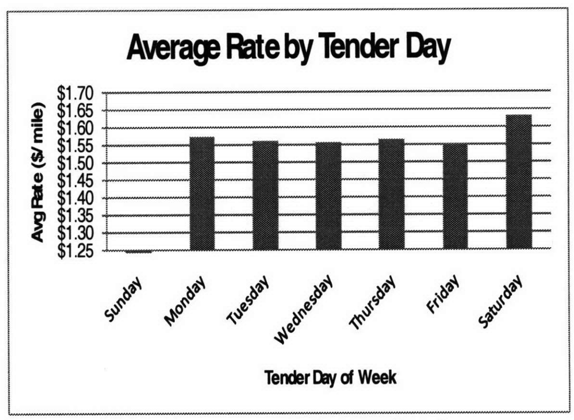

Figure 4.12: Average rate by tender day

A more promising result presented itself at a broad level related to the day on which loads were scheduled to be picked up. Friday stands out as a particularly high volume day for loads and has over 7% more volume than Thursday. Even Saturday and Sunday show activity accounting for over 7% of the volume. Saturday was the day 5.2% of loads were picked up. This trend in the dataset would seem to make sense with warehouses and administrative staffs based on a Monday-Friday work week and measured using performance metrics that conclude at the end of the week. This could possibly explain the surge of activity at the end of the week in order to meet quotas with overtime or increased efficiency and results in additional loads being completed Friday afternoon for Saturday pick up. The dataset also showed a significant difference in average rates per week based on the day of the week for load to be picked up. Not surprisingly, Saturday proved to be an expensive day since shippers paid an average of $47 per load more compared to a load scheduled for pick up on a week day.

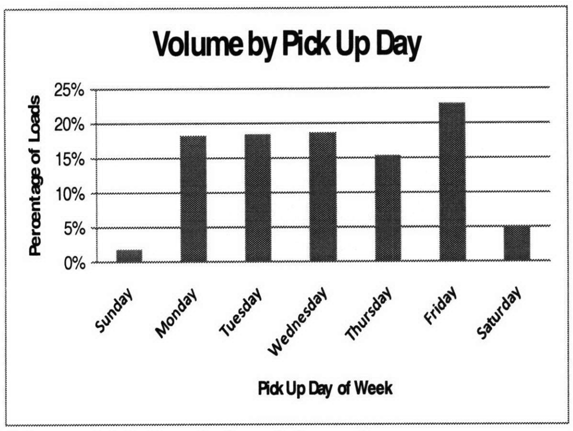

Figure 4.13: Loads by pick up day

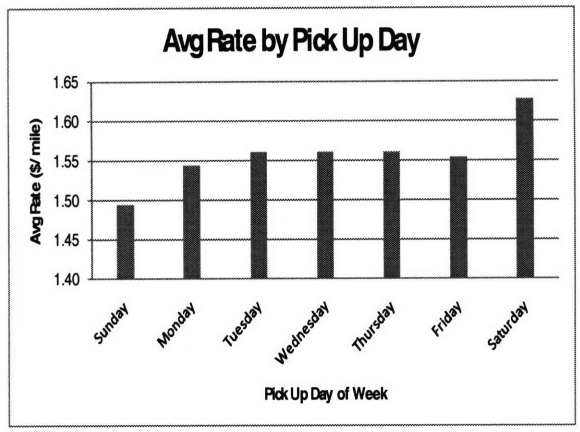

Figure 4.14: Average rates by pick up day

As seen in the data profiling, the transportation model reflects the different costs based on which day of the week a load is tendered. By taking into account the impacting variables, the model provides a better result than the general profile data. Saturday is a very expensive day to call carriers and tender loads, incurring an average penalty of $60 per load. The end of the week tends to be more expensive to tender than the beginning of the week when carriers are trying to plan loads for that week. This intuitively makes sense with load planners working on a weekly basis. The planners then have a natural pattern trying to tender more freight at the end of the week with the result of having to go deeper into the routing guide since more carriers are already committed by Thursday.

The transportation model helps explain the impact of pick up day of the week (see Table 4.6). Just like the case for tendering, Saturday is more expensive than a weekday to pick up freight. In the case of pick up there was also enough volume on Sunday to understand its cost impact on a load and proved to be the most expensive day to have a carrier pick up freight. The other days of the week did not have any meaningful costs or benefits associated with them.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Tender Day	Sunday	N/A			
	Monday	$ (0.38)	0.000	($2.4503)	$1.6918
	Tuesday	Base Case			
	Wednesday	$ (2.07)	0.045	($4.0934)	($0.0430)
	Thursday	$ 4.05	0.000	$1.9939	$6.1001
	Friday	$ 5.94	0.000	$3.7767	$8.0985
	Saturday	$ 49.23	0.000	$29.0772	$69.3746
Pick Up Day	Sunday	$ 16.55	0.000	$11.4502	$21.6450
	Monday	$ 0.61	0.586	($1.5940)	$2.8191
	Tuesday	Base Case			
	Wednesday	$ (1.00)	0.366	($3.1647)	$1.1666
	Thursday	$ 2.29	0.050	($0.0038)	$4.5923
	Friday	$ (1.54)	0.153	($3.6482)	$0.5721
	Saturday	$ 22.98	0.000	$19.6140	$26.3476
Adjusted R^2 Goodness of Fit		0.905			

Table 4.6: Tender and pick up day of week impact in transportation model

4.2.6 Corridor Volume

The corridor volume measures how much volume travels between three digit zip codes on an annual basis. The corridor volume was broken into buckets to better understand the scale of annual volume with daily shipping more than 150 loads, weekly being 52 to 149 loads, monthly being 24 to 51 loads, and annually being less than 24 loads (see Table 4.7). Using the transportation model to quantify the impact of corridor volume we saw that the more loads shipped on a corridor the less expensive the average load becomes. Carriers seem particularly sensitive to corridor volume when the number of annual loads drops below 52 (monthly shipping frequency) and then again when annual loads drop to less than 24 (annual shipping frequency). The difference between the daily and weekly shipping impact was very minor which tells us that even shipping once a week on a corridor allowed carriers to better plan for that volume and increase tender acceptance rates.

Transportation Model					
Variable	Criteria	Impact	P Value	Lower Bound	Upper Bound
Corrider Volume (per year)	Daily >150 loads ($ per load)	Base Case			
	Weekly = 52-149 loads ($ per load)	$ 1.69	0.063	($0.0884)	$3.4636
	Monthly = 24-51 loads ($ per load)	$ 21.27	0.000	$18.7741	$23.7734
	Annual <24 loads ($ per load)	$ 61.51	0.000	$59.1380	$63.8824
Adjusted R^2 Goodness of Fit		0.905			

Table 4.7: Corridor volume impact in transportation model

4.2.7 Factors not Incorporated in Transportation Model

During the analysis there were two variables we reviewed but removed from the transportation model since they did not provide additional insight or allow transportation planners to improve cost performance. The first variable we examined related to consistency of shipping on a given lane with higher consistency being rewarded with lower tender rejection rates since carriers could more easily plan their capacity. We approximated shipper consistency by incorporating a variable for the number of weeks in a year that each customer had loads for the specific corridor. Overall the impact was surprising since the medium volume lanes proved to be much less expensive than the lower volume lanes while the highest volume lanes were substantially more expensive than the medium and lower volume corridors. The corridor consistency variable did not add more statistical significance to the entire model and did not help to better understand how to reduce transportation costs. While other variables in the model allowed shippers to adjust their business policies to reduce costs, it is unlikely that customers could change the consistency with which they ship loads on a given corridor without redesigning their transportation network.

The second variable that we considered was how much impact individual customers had in the dataset. The conclusion was that customers accounted for much of the variability in average load prices. However this was misleading since much of the costs being attributed to the individual customer were actually a result of their behavior patterns. For instance, customers with poor lead times or over reliance on larger carriers were penalized more heavily when we added the customer variable while the overall impact attributed to lead time and carrier size was reduced – this difference between the best and worst customers was substantial at $148 per load. This made sense that individual customers would be rewarded or punished for their transportation policy or differences in rates but makes the model less effective for other companies to review the variables in the model – unless the company happens to mirrors the behavior pattern of one of the customers in the model. With the hope that the model results would have a more universal appeal to transportation planners outside the customers in the dataset we removed this variable from the transportation model.

4.3 Model Evaluation

We evaluated the multiple linear regression-based transportation model in terms of explanatory and predictive performance. The training dataset used to develop the model consisted of accepted tenders from 2007, and the test dataset used to evaluate the model consisted of accepted tenders from 2008. The number of loads and average cost per load for the two data sets are shown below.

Dataset	Year	Number of Loads	Sum of Tender Costs	Average Tender Cost
Training	2007	273,696	$279,491,918	$1021
Test	2008	54,400	$55,092,576	$1013

Table 4.8: Summary of training and test data sets used for transportation model

The full transportation model's predicted tender rates were compared to both the base transportation model as well as a naïve model. The base transportation model was like the full transportation except that its inputs were restricted to miles, origin, destination, and corridor volume. The naïve model, used only for comparison, simply predicts that the cost of each load will equal the average load cost observed in the training data. A residual was calculated for each prediction of each model, as the actual tender cost minus the predicted tender cost. We then calculated the sum of all the residuals and the sum of squares of residuals. We also calculated the total percentage error, by dividing the sum of residuals by the actual sum of tender costs. The results are shown below.

Dataset	Model	Sum of Residuals	Sum of Squares of Residuals	Total Error
Training	Full Transportation	-94,028	7,970,394,169	-0.03%
Test	Full Transportation	-935,742	1,503,781,176	-1.70%
Test	Base Transportation	-743,605	1,512,634,831	-1.35%
Test	Naïve	-459,425	16,898,083,825	-0.83%

Table 4.9: Model evaluation results

We have several observations after examining the results. The sum of residuals is negative for all of the models when used with the test dataset, which is not surprising since the average tender cost for the test dataset was eight dollars less than that of the training dataset. The naïve model has the sum of residuals that is nearest zero, which means that its total predicted cost for all tenders was the most accurate. However, the sums of squares of the residuals show that the naïve model has the worst performance when evaluating individual tenders. The full and base transportation models were both more accurate than the naïve model when predicting individual tender costs. The sum of squares of residuals and total percentage error suggest that the full model is an improvement over the base model.

5 Conclusion

Transportation costs tend to be dominated by just a few key variables – distance, origin, and destination. Even though the other factors, to include lead time, make up less than 1% of predicted long-haul truckload costs they still represent potential cost savings that companies can take advantage of through improved business processes.

5.1 Summary

It becomes clear that transportation planning does not end with network design and carrier rate optimization but should extend into reviewing how business policy affects costs. Specifically, shippers need to review their behavior patterns periodically to identify types of activity that are generally associated with cost penalties – insufficient lead time, weekend tender and pick up, or over reliance on larger carriers when smaller carriers are available. The positive aspect of this research is that all of these variables are controllable and are tied directly to policy or routine. We saw a tremendous disparity between customers with regard to business policies such as lead time. The best customer provided carriers more than five days of lead time on average while the worst averaged less than a day. This difference in lead time behavior between the two customers resulted in an estimated cost penalty of $42 per load. This is a substantial amount when considering that the customer with the worst lead time shipped over 8500 loads last year and paid $323,000 more for the same movement than the customer with the better lead time. This additional cost was roughly equal to 4.1% of the customer's total transportation spend and we expect occurred solely by not contacting the carrier sooner.

5.2 Management Insights

This research would not be complete if it merely identified problems without offering viable solutions to mitigate these impacts. The first step in reducing costs is to optimize the business processes associated with the transportation model variables.

5.2.1 Forecasting Transportation Requirements

In the dataset we observed that tender lead time and truckload costs are correlated. One method of increasing lead time is to forecast transportation needs in advance. Like inventory planning for high turn over items it should be possible to accurately evaluate the high volume lanes and determine a minimum number of expected loads for a given week. Using this demand for a given week, shippers could coordinate well in advance with the preferred carrier on that lane to establish an initial baseline of volume three to four weeks in advance to allow the carrier ample time to plan capacity to meet the shipper's needs. This could be done using the same techniques used to plan inventory with cycle stock and safety stock. The key is to determine the average volume and then use the standard deviation to plan a minimum commitment with the carrier, taking into account the possibility of incurring a penalty for not using a carrier asset. All volume commitments above that level would be tendered as total volume needs for the given week became better understood – ideally still well in advance of the 5 day mark to take advantage of higher acceptance frequency. In the worst case the remaining volume that was not included in the original forecast would be tendered with lead times comparable to the historical level before forecasting.

5.2.2 Reducing Time between Tenders

Another way to reduce the impact of lead time is to speed up the time it takes to tender a load to a carrier or reduce the amount of time a carrier has to decide whether to accept or reject a tender. As shown below, there is a considerable amount of time between tenders.

Routing Guide Depth	Cummulative Time From Initial Tender (hrs)	Time Between Tenders (hrs)
1	0	0
2	10	10
3	14	4
4	19	5
5	25	6
6	30	5
7	32	2
8	37	5
9	40	3
10	41	1

Table 5.1: Cumulative and average time between tenders (hours)

In the introduction we saw that even a highly automated process can still have manual steps. The TMC was required to manually enter a destination appointment time after confirming with the customer receiving the load. This makes sense from a process standpoint but introduces the risk of exceptions. These exceptions could be reduced with a completely automated process where destination customers establish set delivery appointments that can be reserved through EDI. In the dataset overview, we saw that the average time a tender was sent to the second carrier in the routing guide was ten hours while the average tender time by routing guide depth was only five hours. In the case of multiple tenders being rejected this can have a greater impact depending on how long the shipper's policy is regarding tender consideration. From the research there were cases of days of tender time being lost as the load progressed through the routing guide with each subsequent carrier in the guide having a greater chance of rejecting the latest tender with a simultaneous increase in reject rate due to the reduction in lead time at each subsequent carrier.

5.2.3 Strategic Approach to Fleet Assets

From our research we determined that most of the impact to the cost of a load is due to exceptions. The graph below shows the variation of the cost per mile index by lead time. The index is the ratio of a load's cost per mile to the average cost per mile for the specific lane. With additional lead time the variation in cost per mile decreases, making costs more predictable.

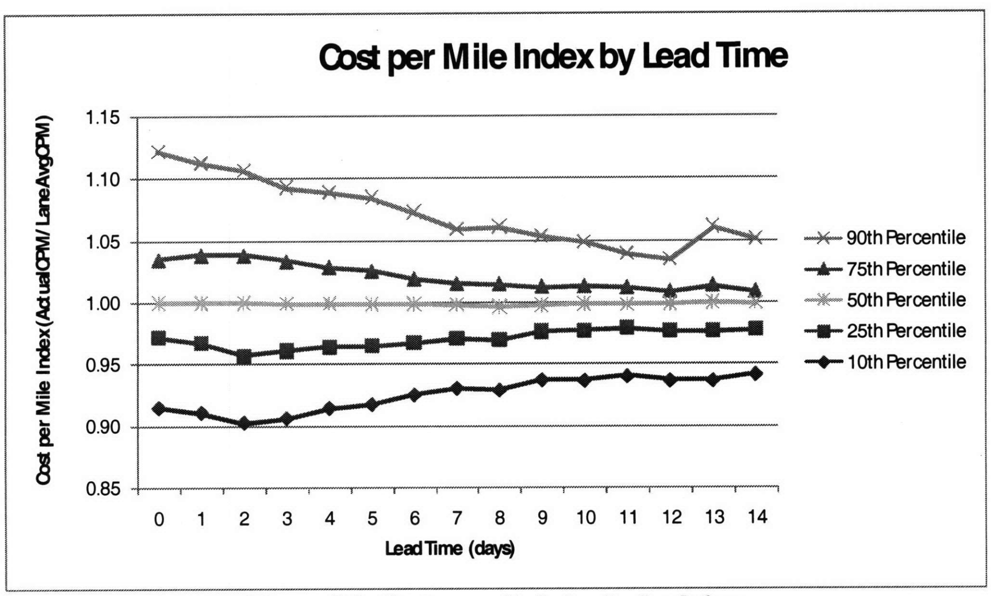

Figure 5.1: Cost per mile index by lead time

In the case of lanes with inconsistent frequency or shorter lead times, companies should explore the use of either private or dedicated fleets in a strategic capacity to improve the worst cost per mile index values, for example those above the 90th percentile.

It would be possible to hold even a limited number of transportation assets in reserve across different corridors and employ them if the initial tenders were rejected by preferred carriers. Rather than continuing deeper into the routing guide at a higher cost and lower service level, the shipper would allocate their own dedicated fleet to handle these loads at a lower cost. Mulqueen (2006) looked at how economies of scale and variability of volume affected the decision to use a dedicated fleet on a lane but based on our research it could potentially make sense to convert to a dedicated fleet much earlier with the additional costs of the transportation model factored into the equation.

Furthermore, our research suggests that managing average costs should be less of a focus than managing the most expensive outliers since they tend to have a greater impact. Companies that maintained even a small, flexible dedicated fleet capable of dealing with transportation exceptions would have the ability to reduce the impact of the worst percentile of loads by handling them internally at a relatively fixed cost. In most cases the internal fleet would be used to handle routine, shorter haul volume until an exception load presented itself, at which point it would be assigned to the difficult load. The routing guide of the routine lane would be used to back fill the internal asset using the standard tender process on the routine lane until it returned from the high priority assignment.

Shippers should also use contractual leverage and rewards to encourage preferred carriers to accept tenders. Since most rates are fixed, transportation cost really becomes a matter of how deep a customer has to go in the routing guide which is tied directly to how often the preferred carrier rejects tenders – 22% of the time on average. From our research we saw how the transportation variables, including tender lead time, affect the average routing guide depth by making loads less attractive to carriers and this in turn translated into increased costs. Regardless of customer behavior it should be possible to have an impact on carrier acceptance behavior. Typically the first two carriers in the routing guide are less expensive than the other carriers. Customers should use volume incentives or acceptance incentives to discourage carriers at the top of the routing guide from rejecting loads. This can be provided as part of a cost savings program where the positive impact of an increased acceptance rate is shared financially on a frequent basis with those carriers. This behavior rewards program needs to be both meaningful financially to change the way the preferred carriers treat customer loads as well as happen at least on a monthly basis or it will not remain in the forefront of the carrier capacity planners' consideration when deciding to accept or reject a tender from the customer. Finally, shippers should publish minimum acceptable tender acceptance rates based on position within the routing guide. These should be used as a benchmark for contract renegotiation and as part of the service level reviews. Carriers that fail to consistently meet these standards of behavior with regard to tender acceptance need to be removed from future contract consideration in a specific lane for a period of time.

5.2.4 Hidden Impact of Metrics on Transportation Costs

From the research we identified the following trends with regard to day of the week:

- Pick up volume is highest on Friday (23-49% more loads than other weekdays)
- Loads tendered Wednesday to Friday are on average $17 more expensive than those tendered on Monday and Tuesday
- Weekend tendering and pick up activity incurs a heavy cost penalty

We believe that the business policies and metrics associated with the typical Monday-to-Friday work week cause a surge in activity in the second half of the week. As a result shippers tender more loads for Friday pick up and exception loads are forced to pick up on the weekend. Assuming consignee flexibility, we believe that changing the weekly performance metrics for warehouses and manufacturers to a Wednesday-to-Tuesday cycle would offset the customer's end-of-week surge, making the loads more attractive to carriers that currently have to plan for a weekly spike in demand every Friday. This would also provide the benefit of unplanned loads occurring on the last day of the week, a Wednesday, with two working days left to handle any exceptions before the more expensive weekend.

Even if customers are unable to change their weekly metric schedule, the transportation planners need to review the frequency and impact of weekend loads. Since carriers prefer to operate Monday to Friday and activity outside of this schedule incurs a penalty, planners can quantify that impact to leadership to evaluate whether the current business plan makes sense to have loads picked up on the weekend. This could be especially interesting if the delivery date could still have been achieved with a Monday pick up at less cost.

Finally, these recommendations are not provided to be an exhaustive list of ways to reduce the impact of transportation variables. The key point of this research was to identify whether lead time has an impact on transportation costs and then quantify that effect. From the research we found that lead time has a definitive impact on transportation cost along with a number of other variables that we used in our transportation model. With an increased understanding of how business policy affects shipping costs we hope that companies will be able not just employ some of our suggestions to optimize lead time and other critical factors but also explore ways to improve coordination with carriers to the benefit of both parties.

5.3 Future Research

Although outside of the scope of our research, it would be interesting to investigate how qualitative aspects of transportation service affect shipper decisions. Based on our research we saw a strong correlation between routing guide position and cost – lowest cost carriers were generally at the top of the routing guide. One would expect that shippers optimize based on cost but must take into account other factors when establishing the routing guide. If service levels of carriers were tracked by lane it would be insightful to see how these different service factors had an impact on the relative depth of the carrier in the routing guide. This analysis would help carriers decide how much they would be willing to spend on each service factor based on the weighted importance of those factors by shippers.

Attempting to predict tenders per load also revealed several areas in which additional work could be done. We think many of the models could be improved with a larger training set. We would also like to try weighting the training set to make higher routing guide depths more common.

Additional research needs to occur with regard to the Morgan Stanley Index. We expected to see a greater impact based on the Hubbard (2007) research that "doubling the market thickness would lead to a 30% increase" in activity outside the normal channels such as using the spot market to contract carriers. We were able to see a strong correlation between the freight index and the depth shippers were forced to go in the routing guide but were unable to find a correlation between that activity and increased price. This was a surprising result that warrants additional investigation. It is possible that one year of data is insufficient based on the fact that carrier rates are typically negotiated annually and these pricing differences would not necessarily be reflected in the short term. The other concern we had with the 2007 freight index was that it was a very steady year compared to the years before. It would be interesting to see if more extreme values of the freight index would cause a greater impact to rates.

Reference List

Anderson, C. K. & Wilson, J. G. (2003). Wait or buy? The strategic consumer: Pricing and profit implications. *Journal of the Operation Research Society, 54*, 299-306.

Bales, W. A. (1993). CEOs'/presidents' perceptions and expectations of the purchasing function [WWW Document]. CAPS Research. URL http://www.capsresearch.org/Publications/pdfs-protected/bales1993.pdf (visited 2007, October)

Bearth, Daniel (2007). Top 100 For Hire Carriers [WWW Document]. Transport Topics. URL http://www.ttnews.com/tt100/2007/index.asp (visited 2007, December).

Berry, M. & Linoff G. (2004). Data mining techniques: For marketing, sales, and customer relationship management, 2nd ed. Hoboken: John Wiley & Sons, Inc.

Caplice, C. (1996). An optimization based bidding process: A new framework for shipper-carrier relationships. Unpublished doctoral dissertation, Massachusetts Institute of Technology.

Caplice, C. (2006). Electronic markets for truckload transportation. Unpublished paper, Massachusetts Institute of Technology.

Caplice, C. & Sheffi, Y. (2006). Combinatorial auctions for truckload transportation. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), *Combinatorial auctions* (Chapter 21). Cambridge and London: MIT Press.

CHAINalytics LLC (2006). Model Based Benchmarking Consortium (MBBC) market report. Unpublished paper.

Harris, F. H. & Peacock, P. (1995). "Hold my place, please": Yield management improves capacity-allocation guesswork. *Marketing Management*. Fall 1995, 34-45

Hubbard, T. (2001). Contractual Form and Market Thickness in Trucking. *The RAND Journal of Economics*. Summer 2001, 369-386.

Hunter, E. (2007, April 23). As meetings thrive, buyers face shrinking lead times [WWW Document]. Business Travel News Online. URL http://www.btnmag.com/businesstravelnews/headlines/frontpage_display.jsp?vnu_content_id=1003574982 (visited 2007, October)

Mulqueen, M. J. (2006). Creating transportation policy in a network that utilizes both contract carriers and an internally managed fleet. Unpublished master dissertation, Massachusetts Institute of Technology.

Morgan Stanley Research North America (2008). Freight Transportation – Morgan Stanley Propriety Truckload Freight Index. April 4, 2008.

Shmueli, G., Nitin, R., & Bruce, P. (2007). *Data Mining for Business Intelligence.* Hoboken: John Wiley & Sons, Inc.

Swenseth, S. R. & Godfrey, M. R. (2002). Incorporating transportation costs into inventory replenishment decisions. *International Journal of Production Economics, 77,* 113-130.

Zhelev, G. (2001). Flexibility in transportation procurement: A real options approach. Unpublished master dissertation, Massachusetts Institute of Technology.

Appendix

A.1 Description of Data Fields

Data Field	Description
LoadNum	Unique load identifier
OriginCity	City where load is picked up
OriginState	State where load is picked up
DestinationCity	City where load is delivered
DestinationState	State where load is delivered
O_Zip	Zip Code where load is picked up
D_Zip	Zip Code where load is delivered
EnteredDate	Date load entered TMC system
Tender _Seq	Actual tender sequence
Rejected	Binary value to indicate if tender was rejected
TenderedDate	Date tender sent to carrier
BookedDate	Date carrier accepted load
Ship Date	Date load ships
PU By	Date and Time load to be picked up
P_Act_Arrived	Date and Time carrier arrived
Total Rate	Total spend on carrier for load
Tender Rate	Contractual rate excluding accessorial charges at tender
Core/NonCore	Carrier in top two position in routing guide
Mode	Type of load
Carrier	Carrier Name
Carrier Home City	Headquarters city of carrier
Carrier Home State	Headquarters state of carrier
Miles	Total miles of load
Customer	Client company
Lead Time	Difference from order entry to pick up by date
Rate per Mile	Total tendered rate divided by total miles

Table A.1: Definition of key data fields

A.2 Additional Data Profiling Results

State	Origin State	Dest State	Region	Origin Region	Dest Region
CT	0%	0%	NE	24.4%	25.5%
DE	0%	0%			
MA	2%	1%			
MD	0%	1%			
ME	0%	1%			
NH	0%	0%			
NJ	3%	2%			
NY	0%	2%			
PA	7%	7%			
RI	0%	0%			
VA	1%	2%			
VT	0%	0%			
IN	1%	1%			
NC	11%	8%			
MI	0%	0%			
IA	1%	3%	NC	23.3%	19.7%
ID	0%	0%			
IL	3%	4%			
KS	0%	1%			
KY	0%	1%			
MN	2%	2%			
NE	0%	1%			
OH	14%	5%			
SD	0%	0%			
WI	3%	2%			
WV	0%	0%			
AK	0%	0%	NW	1.4%	4.3%
MT	0%	0%			
ND	0%	0%			
OR	1%	1%			
WA	1%	2%			
WY	0%	0%			
CO	0%	1%	SC	16.9%	17.2%
LA	0%	2%			
MO	3%	2%			
MS	0%	1%			
NM	0%	0%			
OK	1%	1%			
TX	12%	10%			
AL	3%	3%	SE	19.8%	19.6%
FL	4%	7%			
GA	6%	5%			
AR	1%	1%			
SC	1%	1%			
TN	4%	3%			
AZ	0%	2%	SW	14.2%	13.7%
CA	11%	9%			
HI	0%	0%			
NV	3%	1%			
UT	1%	2%			

Table A.2: Dataset volume by state and region (origin and destination)

Tender Day	Loads	Volume %	Cost	Avg Rate	Avg Seq
Sunday	0	0.00%		0.00	
Monday	64668	19.59%	$ 1,010	1.57	0.36
Tuesday	66935	20.28%	$ 1,007	1.56	0.38
Wednesday	67754	20.52%	$ 1,024	1.55	0.40
Thursday	68488	20.75%	$ 1,031	1.56	0.41
Friday	61857	18.74%	$ 1,022	1.55	0.42
Saturday	414	0.13%	$ 1,487	1.63	0.33

Table A.3: Pick Volumes by day of week

Pick Up Day	Loads	Volume %	Cost	Avg Rate	Avg Seq
Sunday	6175	1.87%	$ 895	1.49	0.51
Monday	59795	18.11%	$ 1,015	1.54	0.34
Tuesday	60398	18.30%	$ 998	1.56	0.36
Wednesday	61231	18.55%	$ 994	1.56	0.36
Thursday	50200	15.21%	$ 977	1.56	0.35
Friday	75031	22.73%	$ 1,082	1.55	0.44
Saturday	17286	5.24%	$ 1,097	1.63	0.61

Table A.4: Pick volumes by day of week

Sequence	Loads	Volume %	Cost	Avg Rate	Avg Lead Time
1	258533	78.3%	$ 1,021	$ 1.54	3.46
2	44378	13.4%	$ 994	$ 1.59	3.46
3	14656	4.4%	$ 1,007	$ 1.62	3.25
4	5961	1.8%	$ 1,058	$ 1.69	2.95
5	2801	0.8%	$ 1,061	$ 1.79	2.68
6	1414	0.4%	$ 1,153	$ 1.86	2.55
7	850	0.3%	$ 1,139	$ 1.89	2.35
8	493	0.1%	$ 1,203	$ 1.97	2.21
9	344	0.1%	$ 1,241	$ 1.98	2.44
10	211	0.1%	$ 1,319	$ 2.07	3.09
11	150	0.0%	$ 1,258	$ 2.15	2.07
12	104	0.0%	$ 1,304	$ 2.24	2.36
13	61	0.0%	$ 1,318	$ 2.15	1.97
14	33	0.0%	$ 1,270	$ 2.19	2.55
15	26	0.0%	$ 1,376	$ 2.16	1.62
16	28	0.0%	$ 1,245	$ 2.11	2.11
17	20	0.0%	$ 1,123	$ 2.16	1.85
18	16	0.0%	$ 1,061	$ 1.85	2.00
19	12	0.0%	$ 1,336	$ 1.94	1.58
20	12	0.0%	$ 1,381	$ 2.08	1.33
21	3	0.0%	$ 846	$ 1.70	0.67

Table A.5: Dataset routing guide depth factors

Lead Time (Days)	Loads	Volume	Cost	Avg Rate	Avg Seq
0	20900	6.33%	$ 1,062	$ 1.65	1.556
1	62689	18.99%	$ 1,080	$ 1.56	1.485
2	63092	19.11%	$ 1,060	$ 1.59	1.392
3	45593	13.81%	$ 1,033	$ 1.56	1.371
4	44707	13.54%	$ 1,043	$ 1.57	1.337
5	32099	9.72%	$ 973	$ 1.56	1.336
6	26239	7.95%	$ 892	$ 1.52	1.304
7	15803	4.79%	$ 891	$ 1.49	1.328
8	8630	2.61%	$ 855	$ 1.41	1.339
9	2598	0.79%	$ 918	$ 1.41	1.294
10	2066	0.63%	$ 968	$ 1.41	1.399
11	1393	0.42%	$ 1,017	$ 1.43	1.395
12	1176	0.36%	$ 958	$ 1.43	1.277
13	786	0.24%	$ 925	$ 1.44	1.282
14	632	0.19%	$ 911	$ 1.41	1.283
15	311	0.09%	$ 948	$ 1.42	1.347
16	244	0.07%	$ 993	$ 1.37	1.455
17	153	0.05%	$ 973	$ 1.34	1.307
18	179	0.05%	$ 958	$ 1.31	1.302
19	149	0.05%	$ 1,037	$ 1.26	1.215
20	102	0.03%	$ 887	$ 1.36	1.186
21	94	0.03%	$ 903	$ 1.29	1.245

Table A.6: Dataset lead time factors

A.3 Additional Details of Models to Predict Tenders per Load

A.3.1 Neural Network Results

Notes:
a) network had 1 hidden layer with 25 nodes
b) error percentage went up slightly with two hidden layers with 25 nodes each, and with two hidden layers with 15 nodes each

Class	# Cases	# Errors	% Error
1	389	9	2.31
2	117	113	96.58
3	34	30	88.24
4	16	9	56.25
5	12	12	100.00
6	7	7	100.00
7	5	5	100.00
8	1	1	100.00
9	2	2	100.00
10	1	1	100.00
13	1	1	100.00
Overall	585	190	32.48

Table A.7: Neural network test dataset error summary

	Predicted Class										
Actual Class	1	2	3	4	5	6	7	8	9	10	13
1	380	3	1	5	0	0	0	0	0	0	0
2	108	4	3	2	0	0	0	0	0	0	0
3	25	2	4	3	0	0	0	0	0	0	0
4	7	0	2	7	0	0	0	0	0	0	0
5	7	0	4	1	0	0	0	0	0	0	0
6	1	0	2	4	0	0	0	0	0	0	0
7	2	2	1	0	0	0	0	0	0	0	0
8	0	0	1	0	0	0	0	0	0	0	0
9	0	0	2	0	0	0	0	0	0	0	0
10	1	0	0	0	0	0	0	0	0	0	0
13	1	0	0	0	0	0	0	0	0	0	0

Table A.8: Neural network test dataset confusion matrix

A.3.2 Discriminant Analysis Test Dataset Scoring Summary

Class	# Cases	# Errors	% Error
1	389	136	34.96
2	117	117	100.00
3	34	19	55.88
4	16	16	100.00
5	12	12	100.00
6	7	7	100.00
7	5	5	100.00
8	1	1	100.00
9	2	2	100.00
10	1	1	100.00
13	1	1	100.00
Overall	585	317	54.19

Table A.9: Discriminant analysis test dataset error report

	Predicted Class										
Actual Class	1	2	3	4	5	6	7	8	9	10	13
1	253	28	51	27	0	0	0	0	0	0	30
2	80	0	25	4	0	0	0	0	0	0	8
3	17	0	15	1	0	0	0	0	0	0	1
4	8	0	6	0	0	0	0	0	0	0	2
5	4	1	6	0	0	0	0	0	0	0	1
6	0	0	7	0	0	0	0	0	0	0	0
7	2	0	3	0	0	0	0	0	0	0	0
8	0	0	1	0	0	0	0	0	0	0	0
9	1	0	1	0	0	0	0	0	0	0	0
10	0	0	1	0	0	0	0	0	0	0	0
13	0	0	1	0	0	0	0	0	0	0	0

Table A.10: Discriminant analysis test dataset confusion matrix

A.3.3 Classification Tree Overview and Results

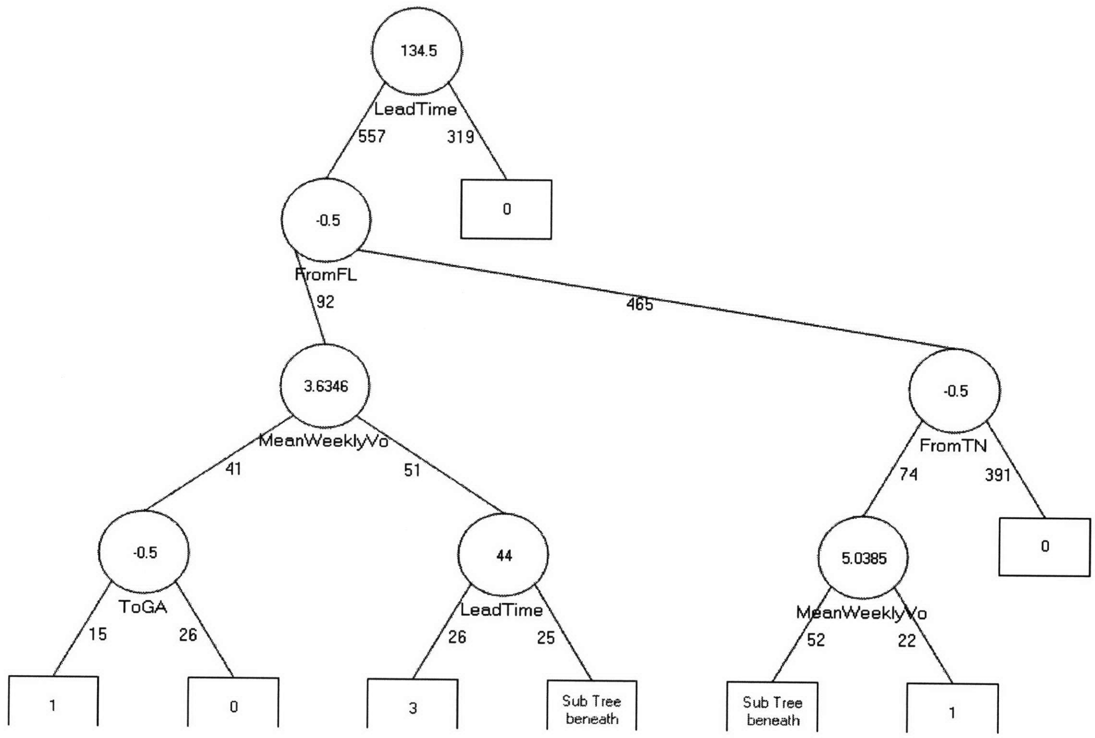

Figure A.1: Top nodes of pruned classification tree

Note: binary variables used in this model are encoded using -1 for true and 0 for false

Class	# Cases	# Errors	% Error
1	389	13	3.34
2	117	107	91.45
3	34	31	91.18
4	16	9	56.25
5	12	12	100.00
6	7	7	100.00
7	5	5	100.00
8	1	1	100.00
9	2	2	100.00
10	1	1	100.00
13	1	1	100.00
Overall	585	189	32.31

Table A.11: Classification tree test dataset error rates

	Predicted Class										
Actual Class	1	2	3	4	5	6	7	8	9	10	13
1	**376**	3	2	8	0	0	0	0	0	0	0
2	100	**10**	1	6	0	0	0	0	0	0	0
3	19	7	**3**	5	0	0	0	0	0	0	0
4	6	2	1	**7**	0	0	0	0	0	0	0
5	6	2	0	4	**0**	0	0	0	0	0	0
6	1	1	0	5	0	**0**	0	0	0	0	0
7	2	2	0	1	0	0	**0**	0	0	0	0
8	0	0	0	1	0	0	0	**0**	0	0	0
9	0	0	0	2	0	0	0	0	**0**	0	0
10	0	1	0	0	0	0	0	0	0	**0**	0
13	0	1	0	0	0	0	0	0	0	0	**0**

Table A.12: Classification tree test dataset confusion matrix (correct predictions in bold)